INTERNATIONAL EDITION · OPÉRATION DINDON · JULY 2026

CORPUS DINDON — INTERNATIONAL EDITION

Thesis, Capture Mechanisms, Legal Dossier, Physical Substrate, and Exit Solutions

◆ THE CENTRAL THESIS

No digital sovereignty without control of matter. This corpus brings together the full body of Opération Dindon studies: the founding thesis, the structural mechanisms of capture — technical, cognitive, human — the complete legal dossier, the physical substrate they conceal, and concrete paths out. One demonstration, in 85 pieces, structurally audited and validated by an independent contradictor.

PIECES	PARTS	PAGES	SINCE
85	11	719	MAY 2026

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Trainer since 2006

🏛️ Nemo supra legem est

Public document · CC BY-NC-SA 4.0 · Opération Dindon

HUMAN

TABLE OF STUDIES

The Teacher's Manifesto — Amine RAITI — BIOGRAPHY P.8
Biography and stance of the author: two decades of teaching and infrastructure work, the origin of this entire corpus.

I — THE THESIS P.9

No Sovereignty Without Matter — Amine RAITI P.10
Founding thesis of the corpus: digital sovereignty is defined by control of matter, not law alone.

II — CAPTURE MECHANISMS P.18

The Infrational Crisis — Amine RAITI P.19
Introduces the neologism "infrational": as inflation dilutes currency, lexical inflation dilutes the meaning of technical words.

The Infrational Loop — Amine RAITI P.26
Documents the loop where the dilution of bare-metal skills drives salaries up, pushes seniors to leave, and worsens the shortage.

The Global Anatomy of Amputation — Amine RAITI P.33
Maps two linked amputations: the infrastructure body's loss of knowledge, and the specific exclusion of women from age twelve.

Anatomy of the Loss — Diagnosis and Reformation — Amine RAITI P.40
Structural diagnosis of the dissolution of historic technical guardrails between development and infrastructure, without judging individuals.

Anatomy of Digital Perdition — Amine RAITI P.50
Models seven causal nodes and five bifurcation points running from primary school to signing a hyperscaler contract.

The Dissolution of the Boundary — Amine RAITI P.66
The merging of dev and infra roles under a single "DevOps" title is an intimate driver of cloud captivity.

The Cloud-Washing — Amine RAITI P.74
Describes the meta-mechanism that repaints dependency as freedom and extraterritoriality as sovereignty, across three registers.

The Free Sample — Amine RAITI P.81
The Free Tier follows the same mechanics as the free sample in retail: create the need before monetising it.

The Free Tier Bait — Amine RAITI P.87
Documents seven mechanisms of the Free Tier, from capturing students to the financial anaesthesia of start-up credits.

The €3,000 VM and the €2,800 SRE — Amine RAITI P.97
A cloud VM costs more per month than the engineer who runs it — and it's the human who leaves first.

The Economy of Silence — Uptime as the Only Metric — Amine RAITI P.105
Provocative thesis: the best SRE is often the one who appears to do nothing, because their infrastructure runs itself.

The Ticket and the Talent — Amine RAITI P.112
Proposes a three-category, six-level nomenclature (SysOps, NetOps, OpInfra) to restore a precise vocabulary to infrastructure.

The Newspeak That Costs Dear — Amine RAITI P.119
Documents, term by term, the measurable effects — balance sheets, org charts, ToS — of cloud newspeak on organisations.

NoOps — Autopsy of a Word That Killed a Profession — Amine RAITI P.131
Autopsy of the one marketing term that explicitly announced a profession's disappearance, adopted without resistance by its own victims.

The Invisible Debt — Amine RAITI P.137
Analyses the world's dependency on unfunded open source (Linux, OpenSSL, curl), exposed by the 2021 Log4Shell flaw.

The Departure of the Last One Who Knows — Amine RAITI P.144
Addresses the attrition and cognitive desertification of the senior SRE, and the tacit knowledge that leaves with them unrecorded.

III — LEGAL DOSSIER P.151

The Architecture of Consent — Amine RAITI P.152
Maps the cognitive biases (Cialdini, behavioural economics) hyperscalers exploit at every layer of the loss of sovereignty.

Terms of Service Under the Microscope — Amine RAITI P.160
Dissects AWS and GCP terms of service clause by clause, showing the corpus's mechanisms literally written into them.

Lock-In by Design: A Six-Layer Taxonomy of Retention Mechanisms in Hyperscaler Cloud Contracts · Amine RAITI P.168
Academic taxonomy of retention mechanisms in AWS, GCP, and Azure contracts, in six layers, based on 47 ToS versions.

Mapping the Stratagems of Vendor Lock-in — Amine RAITI P.180
Documented French case study: an internal diagnostic tool evaluating six lock-in layers across three hyperscalers via a 12-question grid scored out of 24.

The Contractual War Machine — CMA Version · Amine RAITI P.185
Systemic deconstruction of hyperscaler terms addressed to the CMA, FCA, and ICO, analysing clauses against UCTA 1977 and British common law.

The Contractual War Machine — Complete Anatomy · Amine RAITI P.190
Documented French case study: a systemic deconstruction of hyperscaler lock-in into three layers and seven mechanisms, grounded in nineteen source legal analyses.

La Machine de Guerre Contractuelle — Version Trilingue Unifiée · Amine RAITI P.195
Simultaneous French/English/German version aimed at three regulators (France, UK, Germany), each on their own national legal grounds.

Legal Analysis · AWS Service Terms EN vs British Law · Amine RAITI P.211
Legal audit confronting AWS Service Terms with British law, exposing clauses that conflict directly with UCTA 1977 and the Competition Act 1998.

Legal Analysis · Google Cloud Terms of Service vs British Law · Amine RAITI P.218
Frontal analysis of Google Cloud Terms of Service, identifying 9 critical clauses potentially void or unenforceable under current British law.

IV — THE PHYSICAL SUBSTRATE: IRON P.222

Mastery of the Iron — The Final Lock and the Path to Reconquest — Amine RAITI P.223
Closing study of the corpus: the lock on Iron is not a conspiracy but an investment imbalance, reversible through training.

The Digital Iron — Amine RAITI P.230
Extends the logic of physical iron to semiconductors: no chip, no server; no server, no sovereignty.

The Guardians of Iron — Amine RAITI P.237
Documents the hyperscalers' own bare-metal training programmes (STAR Program, Datacenter Academy), figures verified source by source.

The Digital Ground — Amine RAITI P.244
Fills a blind spot in the corpus: the location, ownership, and legal framework of European data centres.

The Taiwan Bottleneck — Amine RAITI P.251
Documents the geopolitical dependency on the Taiwan Strait, the deepest and least visible layer of lock-in.

The Material Black Box — Amine RAITI P.257
Self-hosted bare-metal often runs on non-auditable proprietary firmware — a hardware lock-in that replaces software lock-in.

The FinOps of the Last Gram — Amine RAITI P.261
Proposes direct recovery of precious metals (gold, copper, palladium) from obsolete hardware, rather than standard e-waste channels.

Grade A Refurbished — Amine RAITI P.268
Defends Grade A refurbished server hardware as an economical, ecological, and operationally reliable infrastructure strategy.

The Strategic Refurbished — Amine RAITI P.275
Refurbished hardware isn't just a FinOps choice — it's a geopolitical insurance policy whose value rises with cloud vulnerability.

The Gentle BCP — Amine RAITI P.279
Contrasts business continuity sold as automatic by hyperscalers with the continuity actually practised by sovereign managed-service providers.

The Invisible Plumbing — Amine RAITI P.285
Public-facing explainer: nothing "floats" in the cloud — hard drives, cables, and air-conditioned buildings, no jargon.

أتستبدلون — Cloud Paradox — Amine RAITI P.293
Sarcastic, Quran-referencing reflection on the paradox of trading one's own known server racks for invisible foreign ones.

V — WOMEN IN INFRASTRUCTURE P.299

The Silence of the Rooms — Why Women Are Not in SRE — Amine RAITI P.300
Structural diagnosis of the near-total absence of women in systems and network infrastructure, with concrete levers for change.

The Invisible Amputation — Amine RAITI P.307
Extends The Silence of the Rooms by proposing a structural hypothesis on the link between women's exclusion and hyperscaler interests.

Iron Mixity — Amine RAITI P.314
A macro-economic audit identifying the hidden reservoir of female talent to break the infrastructure shortage through sovereign long apprenticeships.

Maternity as Continuity — Amine RAITI P.330
Beyond entry: proposes maintained connection, full financial protection, and anti-degradation by design for mothers in infrastructure roles.

VI — THE PEOPLE P.337

The Yemeni Dinner — Amine RAITI P.338
Real team cohesion is built through chosen human moments, not corporate seminars — a concrete, costed programme.

The Body First — Amine RAITI P.345
Argues that a €35-per-head hammam-walk-dinner evening builds a team cohesion no corporate seminar can reproduce.

The Pager and the Body — Amine RAITI P.352
Applied risk analysis: what permanent on-call duty does to the body and cognition over time.

The Body That Makes — Amine RAITI P.359
Second stage after Arduino: laser cutters, CNC, 3D printers — code that creates matter, an annual project to build.

The Mentor and the Metal — Amine RAITI P.367
The Maâllem/Metaâllem process of Moroccan craftsmanship is the same channel of tacit transmission that forms the best infrastructure engineers.

The Great Recharge — Amine RAITI P.374
An annual unprogrammed weekend in the mountains produces a cognitive restoration no formatted seminar can match.

The State and the Invisible Body — Amine RAITI P.381
The state protects professions whose disappearance would threaten social life, but not infrastructure engineers — six proposed levers.

The Uniform of the Body — Amine RAITI P.388
The uniform embodies the SysOps/NetOps/OpInfra nomenclature: a trade without a uniform is a trade hard to defend.

VII — EDUCATION & TRAINING P.396

The Blue — What the Diploma Does Not Measure — Amine RAITI P.397
Addresses competent profiles excluded from the market by the diploma/certification filter, with a proposed correction under French law.

From the Well to Head of SRE — Amine RAITI P.405
Eight-episode account of the author's own experience, the origin of every thesis in the corpus — who is writing, and why to believe them.

The Maths Teacher's Bag — Amine RAITI P.411
A 90-second teaching gesture, no reform or permission needed: place a concrete object on the desk after the formula on the board.

Sir, What Is This Good For? — Amine RAITI P.426
Traces the causal chain from the classroom to the boardroom, where learned abstraction becomes a bias against the concrete.

The Digital Craftsmen — Amine RAITI P.434
Proposes the missing fourth layer — legal protection — via an elastic Syntec model inspired by the Compagnons du Devoir.

The Foundation of Iron — Detailed Training Plan 26 Weeks — Amine RAITI P.442
Week-by-week plan of a 910-hour programme: 73% of the time devoted to system, middleware, and network, directly employable skills.

The CPF Misused — Dormant Rights, Foundation of Iron, IDPE — Amine RAITI P.454
Corrects a misallocation — public training funds financing proprietary cloud certifications — by redirecting them to a sovereign curriculum.

VIII — SOLUTIONS: EXIT P.461

Technical Primacy — Amine RAITI P.462
The right of the Principal SRE to the final word on architecture within the budget envelope set by management, not beyond it.

The Gentle Exit — Amine RAITI P.470
Proposes concrete exit paths using tools already available, rather than merely restating the problems already diagnosed.

The Accompanied Exodus — Amine RAITI P.478
Managed services as an operational entry point for organisations wanting to leave the cloud but with no one left who can run bare-metal.

The Sovereign Exodus — Amine RAITI P.487
A liberation architecture across six strata and three time horizons; the school-age bifurcation costs almost nothing, the one at twenty costs everything.

The Sovereign Interface — Amine RAITI P.503
An inverted abstraction layer where the cloud provider adapts to the application, not the reverse — switching provider becomes swapping an adapter.

Cloud Exit Guide · Fiche B · IT — Amine RAITI P.511
Documented French case study: a dependency audit evaluating technical, contractual, and pricing layers per service, accompanied by four concrete technical exit strategies.

Cloud Exit Guide · Fiche C · Finance — Amine RAITI P.516
Documented French case study on Western European markets: a financial matrix detailing the true costs of captivity, from egress fees to double-run migrations.

Cloud Exit Guide · Fiche D · Operational Sequence — Amine RAITI P.521
Documented French case study: an execution timeline detailing five timed steps from D-Day to D+180, including a dashboard for managing vendor blockade scenarios.

Cloud Exit Guide · Fiche E · HR — Amine RAITI P.527
Documented French case study: an analysis of skills and demand lock-in, outlining a three-phase transformation plan alongside local institutional frameworks like CSE and OPCO.

SRE/Ops Cloud Cost Reduction Guide · AWS · GCP · Azure · Global Edition · Amine RAITI P.532
33 technical cost-reduction and clean-up blocks, hyperscaler by hyperscaler, with a reference open-source stack, for a global audience.

VMware to Proxmox VE Migration — Stretched Cluster D1/D2/D3 — Amine RAITI P.544
Five-phase migration plan from VMware/Broadcom lock-in to an open-source stack, across three data centres in a stretched cluster.

VMware to Proxmox VE Migration — Complete Technical Tutorial — Amine RAITI P.554
Thirteen-module operational companion: real commands, configurations, and checklists to execute the Proxmox migration described in the architecture study.

MariaDB Without Trembling — Amine RAITI P.571
Complete technical tutorial for production MariaDB replication: SSH streaming, binlog, GTID, and post-deployment verification.

IX — LLM & ARTIFICIAL INTELLIGENCE P.578

From Amplifier to Substitute · Individual Cognitive Lock-In · Generative AI · Amine RAITI P.579
Theorises the shift from an amplifying use of generative AI to a dependency that atrophies the developer's autonomous reasoning.

The Replacement That Reveals — Amine RAITI P.587
Inspired by a LinkedIn post: AI doesn't replace humans, it reveals their rarity as generated content multiplies.

Deamputation at the Source — Amine RAITI P.594
The amputation has an origin in Year 8 classrooms: a €35 Arduino kit and €10 a month are enough to interrupt it.

The Anti-Amputation Foundation — Amine RAITI P.602
A free teaching sequence, requiring no official approval or extra budget, for teachers who want to act without waiting for reform.

The Importance of Partying Often at Work — Comic Annex — Amine RAITI P.608
A live demonstration of the Claude/Gemini/human-arbiter process itself — a study of AI collaboration in action, not a cultural aside.

The Thought Machine — Amine RAITI P.621
AI is neither magic nor consciousness: a numerically controlled machine whose raw material is human thought, to be calibrated.

The Mystical Machine — Amine RAITI P.630
"Artificial Intelligence" is a false friend like "Sovereign Cloud": a CNC machine and an LLM differ only in their layer of work.

X — CULTURAL DIWÂN P.638

Digital Diwân of the Craftsman — EN — Amine RAITI P.639
Seven Mu'allaqat (suspended poems) reprising the corpus's theses — sovereignty, freedom, transmission — in the pre-Islamic Arabic poetic tradition.

The Diwan of the Ten Fires - Amine Raiti P.647
A poetic indictment born of an oral warning received for presenting figures at a management committee — fourteen episodes of nassiha.

The Hand Over Nations — Amine RAITI P.671
A finding of effect, not intent: entire nations' critical infrastructure sits under foreign jurisdiction and law, measurably so.

The Legend of Dindon — Amine RAITI — Operation Dindon P.679
A collection of fourteen episodes published on LinkedIn, the saga of a digital farmer from cloud hype back to bare-metal.

Amine RAITI — Letter of a Witness P.698
An open letter to Pope Leo XIV on hyperscaler practices, written as a duty of universal witness beyond any one faith.

XI — SYNTHESIS P.702

Original Theories and Concepts — Amine RAITI · Opération Dindon P.703
An inventory of the original concepts and neologisms coined by the corpus, so their reach is not lost in the volume.

Opération Dindon — Magisterial Map — Amine RAITI P.708
The map of the entire corpus: eight thematic arcs, one central thesis, one ultimatum, one date — how it all fits together.

OVERVIEW

WHAT THIS CORPUS DEMONSTRATES

◆ THE CENTRAL THESIS

No digital sovereignty without control of matter. Opération Dindon documents, across 85 pieces in 11 parts, how hyperscalers capture organisations — legally, cognitively, economically, and humanly — that entrust them with their infrastructure, and proposes concrete paths out, from the school curriculum to the boardroom.

READING STRUCTURE

- **I** THE THESIS
- **II** CAPTURE MECHANISMS
- **III** LEGAL DOSSIER
- **IV** THE PHYSICAL SUBSTRATE: IRON
- **V** WOMEN IN INFRASTRUCTURE
- **VI** THE PEOPLE
- **VII** EDUCATION & TRAINING
- **VIII** SOLUTIONS: EXIT
- **IX** LLM & ARTIFICIAL INTELLIGENCE
- **X** CULTURAL DIWÂN
- **XI** SYNTHESIS

◆ PRODUCTION AUDIT TRAIL NOTE

During final compilation of this English edition, a duplicate production artefact of Study 65 was detected in the raw manifest (previously appearing under the alternate translation "The Hidden Reservoir"). A strict comparison (99.5% textual similarity) confirmed the duplication. The artefact has been removed. The canonical title **"Iron Mixity — The Hidden Reservoir of Sovereign Digital"** has been retained, preserving the rejected title as a subtitle and maintaining 1:1 indexing parity with the original French volume. Six studies (Fiche B, C, D, E, the Vendor Lock-in Stratagem Map, and the Contractual War Machine's neutral edition) are presented as documented French case studies, translated in full with their original French/EU legal and financial references preserved, rather than adapted with unsourced figures.

THE TEACHER'S MANIFESTO

I speak to you as a teacher — not as an expert.

I started teaching in **2006**. I have trained hundreds of learners since — and I am not finished. Engineering school students, technicians in career transition, professionals in continuing education. Men and women. Juniors of 18 and seniors of 60. Profiles without a diploma and managers with twenty years of experience.

I have taught **Linux client/server, Windows client/server, Cisco networking, OSI model, WiFi, electricity, number systems, microcomputing, programmable logic controllers, Kubernetes**. Everything that is today in the Foundation of Iron — I taught it in the classroom, face to face with learners, with machines under their fingers.

◆ WHY I ALLOW MYSELF TO SPEAK

When I publish on training, the CPF, the IDPE, the infrationary crisis, digital sovereignty — I am not theorising. I am speaking to people I trained. Some of them are reading these lines from a position as CIO, operations manager, or SRE architect at a major company. I watched them learn to cable their first switch. I know what they are capable of. And I know what the market is trying to make them forget.

When I address hyperscalers, I am not speaking on behalf of a think tank. I am speaking on behalf of my students — those whose skills their certifications seek to render obsolete, those whom the title "DevOps" seeks to dissolve, those whom the shortage of bare-metal training seeks to isolate in a market they can no longer leave.

◆ AS A TEACHER

I speak to those I trained with the same frankness as in the classroom — no condescension, no demagoguery. The truth about the market, about competences, about what titles do not measure.

◆ AS A DEFENDER

I speak to hyperscalers with the same frankness as to my students — no personal hostility, with a clear limit. What you are doing to bare-metal competence, I name it. And I will continue.

SUBJECTS TAUGHT · 2006 — TODAY

Linux client/server · Windows client/server · Cisco Networking · OSI Model · WiFi · Electricity · Number Systems · Microcomputing · PLC & Automation · Kubernetes

AMINE RAITI

Infrastructure Architect & SRE · Teaching since 2006 · Electricity to Kubernetes · All audiences

🏛️ NEMO SUPRA LEGEM EST · CC BY-NC-SA 4.0 · OPÉRATION DINDON

OPÉRATION DINDON · INTERNATIONAL EDITION

PART I

THE THESIS

1 piece

HUMAN

FOUNDING STUDY · OPÉRATION DINDON · JUNE 2026

NO SOVEREIGNTY WITHOUT MATTER

Philosophy of Bare-Metal and the Condition of Digital Sovereignty

RATIO

◆ POSITION OF THIS STUDY IN THE CORPUS

This study is the founding thesis of which all other Opération Dindon corpus studies are extensions. It poses the philosophical question the corpus had never explicitly formulated: what is sovereignty? What is the relationship between sovereignty and matter? Why does purely extraterritorial cloud represent a dissolution of sovereignty — not merely a technical risk? From Bodin to Kubernetes — what controlling infrastructure really means.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · SOVEREIGNTY — DEFINITION AND CONDITION

THE HUMAN IS SOVEREIGN WHEN THEY CAN ACT ON MATTER

Sovereignty is not an abstract legal attribute. It is a concrete capacity — the capacity to act, decide, and intervene on reality without depending on a third party to exercise that capacity. This definition precedes and conditions all its legal, political and organisational forms.

◆ FROM BODIN TO SCHMITT — SOVEREIGNTY AS MASTERY OF TERRITORY

Jean Bodin, in 1576, defined sovereignty as "the absolute and perpetual power of a Republic". But behind the legal formula lies a physical reality: this power only exists if it is exercised over a real territory, with real subjects, over accessible matter. Carl Schmitt clarifies four centuries later: "Sovereign is he who decides on the exception." Deciding the exception presupposes that it can be imposed — which presupposes control of the substrate on which that decision applies.

Thomas Hobbes, in Leviathan, links sovereignty to protection: the sovereign protects their subjects in exchange for obedience. But protecting presupposes the capacity to intervene — to deploy force on a territory. A sovereign who does not control their territory cannot protect their subjects. They are no longer sovereign — they are a tenant.

◆ THE METAPHOR OF THE FARMER AND THE TENANT

The landowner-farmer is sovereign over their production. They decide what to plant, when to harvest, how to manage their soil's resources. If the harvest fails, it is their responsibility — their problem to solve. The tenant who rents land from a foreign owner who can reclaim it at any time, modify lease conditions unilaterally, or forbid certain crops — this tenant is not sovereign. They are precarious. They produce on soil that does not belong to them, under rules they did not set, with the permanent risk of termination.

This distinction — owner vs tenant, sovereign vs precarious — is exactly the one that separates sovereign bare-metal infrastructure from extraterritorial cloud.

◆ NASSIHA — SOVEREIGNTY IS NOT A REFUSAL OF THE MARKET

This study does not defend digital autarky. It defends the capacity to choose — to negotiate from a position of strength, to leave if conditions become unacceptable, to build outside the cloud if the situation requires it. Sovereignty is the exit option. Without this option, there is no market — there is captivity.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · BARE-METAL AS DIGITAL TERRITORY

THE CABLE IS THE SOIL. THE RACK IS THE LAND. THE SERVER IS THE FOUNDRY.

The territorial metaphor is not a metaphor. It is an exact functional description. Territory is what conditions the existence of everything deployed on it. For digital infrastructure, this physical substrate — servers, cables, racks, datacentres — is territory in the literal sense.

◆ WHAT BARE-METAL IS — A FUNCTIONAL DEFINITION

Bare-metal is the physical substrate of digital infrastructure. It is the layer that cannot be further abstracted — below bare-metal lies electricity, materials physics, semiconductor chemistry. Bare-metal is the physical server, the disks, the network cards, the cables, the switches, the routers, the redundant power supplies. It is what cloud abstracts behind an API and a smiling console.

Mastering bare-metal means mastering digital territory. It means knowing where data is physically stored, in which datacentre, in which country, under which jurisdiction, on which hardware, administered by whom. Without this mastery, one does not know what one controls — and what one does not know, one cannot defend.

◆ DIGITAL SOIL — WHAT "THE DIGITAL GROUND" DID NOT STATE EXPLICITLY

The study "The Digital Ground" documented the geography of datacentres, the AMS-FRA-DUB concentration, Interxion's acquisition by Digital Realty for $8.4 billion. It did not pose the philosophical question explicitly: why is datacentre geography a sovereignty question and not merely a latency question?

The answer lies in the territoriality of law. The law applicable to a datacentre is the law of the country where it is physically located — and the law of the country of residence of the company operating it. An AWS datacentre in Ireland is simultaneously subject to Irish law, European law, and American law via the CLOUD Act — because AWS is an American company. The data is physically in Europe. It is legally accessible from the United States. This is variable-geometry sovereignty — and variable-geometry sovereignty is nominal sovereignty.

RAITI

◆ THE CHIP AS INGOT — WHAT "THE DIGITAL IRON" ESTABLISHED

"The Digital Iron" documented that the chip is the ingot of the 21st century. This formula is not a metaphor: it describes a geopolitical reality. Controlling chip production means controlling the substrate of all global digital infrastructure — exactly as controlling gold mines controlled economies in the Middle Ages. TSMC in Taiwan, ASML in the Netherlands, Chinese rare earths — the territory of bare-metal begins in the mines and foundries, not in the datacentres.

Amine RAITI · CC BY-NC-SA 4.0

WHAT PURE CLOUD EXTRACTS FROM CONTROL

Cloud is not inherently bad. It is problematic when it is pure — meaning when it is the only option, when the exit option has disappeared, and when the physical substrate is out of reach and out of the jurisdiction of the person depending on it. In this configuration, cloud is not a service — it is organised dispossession.

◆ WHAT CLOUD EXTRACTS — THE FOUR DIMENSIONS

Physical location: where is my data? On which server, in which datacentre, in which city, in which country? The AWS console says "eu-west-1" — Paris or Ireland. It does not say the building, the rack, the server, the disk. Abstraction is total. The user does not know where their data physically is — and they are not entitled to know.

Applicable jurisdiction: the CLOUD Act of 2018 authorises the American government to access data stored on servers of American companies, regardless of their geographic location. Data hosted at AWS Frankfurt is legally accessible by the US Department of Justice without notification to the European user. This is legally framed and assumed extraterritoriality.

Unilateral termination capacity: AWS, Azure and GCP can suspend an account within hours for terms of service violation. No prior recourse, no obligation of sufficient notice to migrate. A state or critical organisation depending 100% on a hyperscaler cloud can be cut off from its infrastructure by a private American decision.

Pricing power: egress fees, contractual commitments, opaque service pricing — once dependent, the user has no negotiating power. The provider sets prices. The user pays.

◆ CLOUD AS TERRITORY RENTED WITH UNILATERAL TERMINATION CLAUSE

The tenant metaphor applies exactly. The organisation that migrates its entire infrastructure to AWS rents foreign digital territory. The rental contract can be modified unilaterally. The lease can be terminated without sufficient notice. The applicable law is not the tenant's law. And the tenant cannot inspect the subsoil of their own land. This is not sovereignty — it is digital precarity.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE THREE CONCRETE VULNERABILITIES

LEGAL · ECONOMIC · OPERATIONAL

◆ VULNERABILITY 1 — LEGAL: THE CLOUD ACT AND EXTRATERRITORIALITY

The Clarifying Lawful Overseas Use of Data Act (CLOUD Act, 2018) authorises American authorities to access data controlled by companies subject to American law, regardless of physical storage location. AWS, Azure, GCP are American companies. Their data — including that stored in European datacentres — is potentially accessible to the US DOJ via an American judicial warrant, without mandatory notification to the European user. This extraterritoriality directly conflicts with European GDPR. The organisation hosting personal data of European citizens at an American hyperscaler is structurally in jurisdictional conflict — a conflict it cannot resolve by its own decision alone.

◆ VULNERABILITY 2 — ECONOMIC: LOCK-IN AS CAPTURE

"Anatomy of the Loss" documented the mechanisms of technical lock-in — proprietary formats, services without sovereign equivalents, certifications that chain teams to platforms. The economic consequence is documented in "The Digital Ground": organisations that migrated to cloud for economies of scale discover 18 to 36 months later that their costs have increased. Egress fees — data exit charges — are the economic mechanism that makes lock-in permanent. One enters cloud for free. One exits paying. This is the pricing structure of a territory one cannot leave.

◆ VULNERABILITY 3 — OPERATIONAL: DEPENDENCE ON A THIRD PARTY FOR CRITICAL FUNCTIONS

In July 2021, an Akamai outage simultaneously took offline websites of several governments, central banks and public health services. In December 2021, an AWS us-east-1 outage paralysed dozens of critical services for several hours. These incidents are not accidents — they are the structural demonstration that concentrating infrastructure with a few providers creates systemic failure points. A state depending 100% on AWS for its critical digital services entrusts its operational continuity to a private American company that has no public service obligation toward it.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · REBUILDING SOVEREIGNTY

THREE NECESSARY CONDITIONS — NONE SUFFICIENT ALONE

Digital sovereignty is not decreed. It is built — through human, technical and political investments that articulate and reinforce each other. Three conditions are necessary. Each is insufficient without the other two.

◆ CONDITION 1 — THE ENGINEERS WHO KNOW

Digital sovereignty rests on a professional body capable of building, operating and maintaining infrastructure independent of hyperscalers. This body is the bare-metal SRE, the systems administrator, the network engineer documented in "The State and the Invisible Body". Without this body, any sovereignty policy remains nominal — one may decree preference for European actors, but if nobody knows how to operate OVHcloud's datacentres, the policy is empty. "The Departure of the Last One Who Knows" documented the cost of this knowledge disappearing. "The Mentor and the Metal" documented the only mechanism that allows it to be transmitted. "The Foundation of Iron" is the concrete proposal to rebuild it.

◆ CONDITION 2 — THE TOOLS THAT ENABLE

Sovereign European actors exist — OVHcloud, Ecritel, Scaleway, Hetzner, Infomaniak, Clever Cloud. They have datacentres, offerings, teams. They cannot absorb the potential demand for want of qualified personnel (condition 1) and for want of a preference policy that would orient markets toward them (condition 3). Sovereign infrastructure exists in its elementary bricks. It does not exist as a credible alternative at scale — because the other two conditions are not met.

◆ CONDITION 3 — THE POLITICAL DECISION

The state that advocates digital sovereignty in its speeches and hosts its own data at AWS practises a contradictory injunction. The necessary political decision includes: European preference in digital public procurement, mandatory certifications for public administration providers, infrastructure training funding (Foundation of Iron via France Travail), RNCP protected title for infrastructure roles, annual report on the state of the body. These six levers were documented in "The State and the Invisible Body". Without political decision, conditions 1 and 2 remain unactivated potential.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · GRADUAL SOVEREIGNTY — A REALISTIC PROPOSAL

NEITHER ALL-CLOUD NOR ALL-BARE-METAL — THE SOVEREIGNTY ARCHITECTURE

Absolute digital sovereignty is a horizon, not a near-term achievable state. What is immediately achievable is a gradual sovereignty architecture — an approach that distinguishes data and systems by criticality and sensitivity, allocating each category to the infrastructure that guarantees the appropriate level of control.

◆ THE GRADUAL SOVEREIGNTY ARCHITECTURE

Tier 1 — Critical infrastructure: health data, judicial data, tax data, defence systems, energy networks, critical telecoms. On sovereign physical infrastructure, in France or the EU, operated by certified, auditable actors, subject to French and European law only. Zero extraterritorial cloud.

Tier 2 — Sensitive infrastructure: citizens' personal data, administrative systems, non-critical health systems, non-critical judicial systems. On European sovereign cloud (SecNumCloud, German C5) or hybrid infrastructure with portability clauses and annual third-party audit.

Tier 3 — Non-critical infrastructure: public websites, non-sensitive communications, productivity tools. Free cloud with contractual real portability clauses and egress fee caps. The Opération Dindon ultimatum is precisely this demand: resiliable commits, reasonable egress fees, real portability, CLOUD Act limited to its territory.

◆ WHAT THIS ARCHITECTURE CONCRETELY REQUIRES

— Engineers capable of operating Tier 1 (condition 1 — the bare-metal SRE body)
— European actors capable of absorbing Tier 2 (condition 2 — OVHcloud, Ecritel, Scaleway)
— Public policy that funds Tiers 1 and 2 and regulates Tier 3 (condition 3 — preference in public procurement)
— And an engineering culture that understands why these distinctions exist — what the Foundation of Iron has sought to build from the beginning.

◆ NASSIHA — SOVEREIGNTY IS AN OPTION, NOT A PRISON

This study does not defend closure. It defends the capacity to open and close at will. A sovereign state can choose to work with foreign partners — because it has the competences to do without them if needed. This exit option defines sovereignty. Without it, there is no partnership — there is dependence.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE CONDENSED THESIS — WHAT THE CORPUS ALWAYS SAID

FROM BODIN TO KUBERNETES — THE SAME TRUTH

The Opération Dindon corpus has produced more than twenty structural studies. Re-reading these studies with the perspective of this founding thesis, one realises they all say the same thing — from different angles, with different data, for different audiences. This same thing is the thesis of this study.

◆ THE THESIS IN ONE SENTENCE

No digital sovereignty without mastery of bare-metal.

And no mastery of bare-metal without the engineers who understand it,

without the organisations that operate it,

and without the states that protect them.

◆ WHAT EACH CORPUS STUDY SAID ABOUT THIS THESIS

Anatomy of the Loss → how mastery of bare-metal is lost through cognitive capture.

The Digital Iron → that the bare-metal substrate begins in chip mines and foundries.

The Digital Ground → that datacentre geography is a question of territorial sovereignty.

The Infrational Crisis → that sovereignty is first lost in words before being lost in competences.

The State and the Invisible Body → that the state does not protect the professional body that is the condition of its own sovereignty.

The Infrational Loop → how the market destroys this body without anyone having decided it.

The Mentor and the Metal → that knowledge transmission is the only mechanism that keeps this body alive.

The Foundation of Iron → that reconstruction begins before secondary school, in classrooms and workshops.

RATIO

The human is sovereign when they can act on matter.
Purely extraterritorial cloud removes this capacity.
Bare-metal restores it.
Everything else is consequence.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART II

CAPTURE MECHANISMS

16 pieces

HUMAN

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE INFRATIONAL CRISIS

When words lose their weight

RATIO

◆ **CONTEXT OF THE STUDY**

This study introduces a neologism — *Infrational* — to name a documented but unnamed phenomenon. Just as monetary inflation dilutes the value of currency, the inflation of technical words dilutes the value of concepts. When "infrastructure" can mean a VM created by Terraform, when "DevOps" can mean a job halfway between two distinct professions, words stop protecting the realities they once named. And when words give way, skills follow.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE NEOLOGISM

INFRATIONAL — WHY THIS WORD WAS NECESSARY

A neologism is justified only if it names something that had no name yet. *Infrational* is justified. It names a precise phenomenon: the progressive devaluation of technical words through overuse, to the point where those words no longer protect the realities they named. Exactly like monetary inflation — when too much currency circulates for the same amount of goods, currency loses its value. When too many distinct realities are named by the same word, the word loses its value.

◆ ETYMOLOGY AS A STARTING POINT

Infra comes from Latin — below. *Structure* — the foundation, what supports everything else. Infrastructure is what lies below and carries. This etymology is not a grammatical detail — it is a functional definition. Infrastructure is the low layer, physical, tangible, that conditions the existence of every layer above it. It cannot be code. It can be *managed* by code, *described* by code, *documented* by code. But it remains metal, cable, electricity, heat and disks that fail.

◆ WHY THE LOSS OF MEANING IS AN OPERATIONAL PROBLEM

This is not a semantic debate. When the word "infrastructure" loses its physical dimension in the everyday language of technical teams, engineers trained in that linguistic environment do not learn to think in physical layers. They learn to think in cloud resources. When the real failure comes — the one that is in the cable, in the network card, in the redundant power supply that no longer is — nobody knows where to look. Because nobody learned that this is where you look.

◆ NASSIHA — WHAT THIS STUDY DOES NOT CLAIM

This study does not claim that cloud is bad, that Terraform is useless, or that DevOps is a mistake. It claims that word choices have real consequences on the skills that develop or do not develop. And that some of these lexical choices were not innocent.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

METAL HAS NEVER BEEN CODE

"Infrastructure as Code" has become one of the central concepts of the DevOps industry. The original idea — managing system configuration with the same rigorous practices as software code, versioning, testing, peer review — was intellectually sound. The problem is not the intention. The problem is the word.

The word "Infrastructure" in "Infrastructure as Code" does not refer to infrastructure. It refers to the configuration of virtual resources — VMs, security groups, load balancers, S3 buckets. These resources exist as parameters in a hyperscaler's console. The real infrastructure — the physical servers on which these resources run, the racks in which those servers are installed, the cables connecting them, the electrical supply powering them, the cooling keeping them cool — is not code. It has never been code. It cannot be code.

◆ THE ANALOGY THAT EXPOSES THE ABSURDITY

AirportAsCode. MotorwayAsCode. NuclearPlantAsCode.

Nobody would dare call "code" the concrete of runways, the steel of bridges or the turbines of reactors. Nobody would claim that a YAML configuration file "is" an airport. Yet this is exactly what "Infrastructure as Code" normalises for computing systems. A Terraform file that creates a VM is not infrastructure. It is code that sends an API request to a hyperscaler to allocate resources on its physical infrastructure — which you will never see, whose state you do not know, and which you do not control.

◆ WHAT IT SHOULD HAVE BEEN CALLED

"Configuration as Code" — precise, exact, unambiguous. The declarative, versioned description of cloud resource configuration. This word would have described the same practice without claiming that code can be physical infrastructure. It was not chosen. "Infrastructure as Code" was preferred — because it gives the illusion that mastering Terraform means mastering infrastructure. And this illusion serves the interests of those who sell access to the real infrastructure.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

DILUTION BY SUFFIX

DevOps was born from a real observation and a legitimate intention. In 2009, Patrick Debois and his contemporaries observed a structural wall between development teams — who deliver code — and operations teams — who maintain systems. This wall produced slow deployment cycles, frequent production failures, and an absence of shared responsibility for system availability. DevOps was a philosophy to tear down this wall through culture, practices and tools. The intention was sound.

But a sound philosophy with a poor name produces poor effects at scale. "DevOps" fused two distinct professions under one word — and this lexical choice opened the door to exactly what it sought to prevent.

◆ THE DANGER OF THE WORD "DEVOPS" — EVEN WELL-INTENTIONED

A developer and a systems engineer have fundamentally different skills, acquired through different paths, exercised in different contexts. Their collaboration is valuable. Their fusion under a single title does not create a more complete profile — it creates a blurrier one. A "DevOps engineer" who excels at CI/CD may know nothing about managing a network incident at 3am. A "DevOps engineer" who masters system diagnosis may write mediocre code. Fusing the names did not fuse the skills. It diluted the evaluation criteria for both.

◆ THE XXXOPS PROLIFERATION — RUNAWAY INFLATION

DevOps was followed by FinOps, SecOps, MLOps, DataOps, GitOps, CloudOps, PlatformOps. Each new "Ops" suffix promises the same thing: the fusion of two worlds that worked in silos. Each new suffix dilutes the meaning of the word "operations" a little further. When everything is called "Ops", the word no longer designates anything precise. And when "operations" no longer designates anything precise, real operational skills — those that maintain systems in production at night — become invisible in frameworks, job descriptions and training plans.

◆ NASSIHA — THE RIGHT QUESTION IS NOT "DEV OR OPS?"

The right question is: what skills are needed to keep this system in production at 3am, and who holds them in the team? This is not a title question. It is an operational reality question. And it deserves a precise answer — not a generic title that reassures recruiters without committing anyone to anything.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE CONSEQUENCES

WHAT WORD INFLATION HAS PRODUCED

Technical word inflation is not a cultural phenomenon without consequences. It has produced three concrete, measurable results in the digital infrastructure market.

◆ CONSEQUENCE 1 — A GENERATION WITHOUT THE PHYSICAL LAYER

Thousands of engineers trained since 2015 learned to "provision infrastructure" without ever touching a physical server, cabling a switch, or diagnosing a network fault on bare metal. They learned to manipulate abstractions — cloud resources, Terraform files, CI/CD pipelines. These skills are real and useful. But they stop at the API layer. Below the API layer there is a physical infrastructure that nobody taught them existed — because the vocabulary gave them the illusion that it did not exist, or did not need to be known.

◆ CONSEQUENCE 2 — CLOUD BUDGETS EXPLODING WITHOUT ANYONE UNDERSTANDING WHY

Organisations that migrated "to the cloud to achieve economies of scale" typically discover, 18 to 36 months after migration, that their infrastructure costs have increased rather than decreased. Part of this increase is structural — cloud pricing models are designed to grow with usage. But another part is avoidable and stems directly from ignorance of the physical layer: overprovisioning because nobody knows how to size, egress fees ignored because nobody understood there were exit charges, cloud-native architectures replicating on-premises patterns without optimising them.

◆ CONSEQUENCE 3 — A MARKET SHORT OF FULL-STACK INFRASTRUCTURE ENGINEERS

The market is critically short of engineers capable of mastering infrastructure from the physical layer to the application layer — from electricity to service. This profile existed. It was called systems administrator, infrastructure engineer, SRE. It was progressively replaced in job descriptions by "DevOps engineers" and "cloud engineers" whose skills stop at the API layer. The complete profile did not disappear because the competence became impossible to acquire. It disappeared because the vocabulary stopped valuing it — and the market follows vocabulary.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE STRATEGY BEHIND THE INFLATION

THE INFLATION WAS NOT ACCIDENTAL

Infrational inflation is not the result of collective carelessness in word choice. It is the result of a deliberate strategy — not necessarily coordinated, but consistent in its effects — by the actors who have the most to gain from the dilution of infrastructure skills.

◆ WHO BENEFITS FROM THE "INFRASTRUCTURE AS CODE" ILLUSION?

AWS, Azure and GCP benefit directly from the belief that mastering Terraform or CloudFormation equates to mastering infrastructure. If engineers believe that infrastructure is code, they do not need to understand what lies beneath the code. And if they do not need to understand what lies beneath the code, they never have reason to ask whether what lies beneath the code might belong to someone other than a hyperscaler. The AWS console that "provisions infra" is a rental interface, not a mastery tool. The vocabulary erases the distinction.

◆ WHO BENEFITS FROM THE "DEVOPS" DILUTION?

Two actors benefit from profession dilution. Employers who replace two specialists with one cheaper generalist — the "DevOps" title justifies paying less for two distinct expertises under a single salary. And tool publishers who sell "DevOps" platforms claiming to solve through tooling what only culture and competence can solve. The word "DevOps" has become a market. Like all markets, it is shaped by the actors who benefit from it.

◆ THE LINK WITH "ANATOMY OF THE LOSS"

The first study of the Opération Dindon corpus documented the dissolution of infrastructure competence and cognitive capture by hyperscalers. The Infrational crisis is the linguistic mechanism that made this dissolution possible. You cannot lose what you cannot name. You cannot name what you have lost the definition of. Word inflation preceded and conditioned the dissolution of skills. This is not a coincidence.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

TAKING BACK CONTROL OF WORDS

Taking back control of words is not an exercise in nostalgia or lexical purism. It is an operational necessity. The words we use to describe our profession shape the skills we develop, the hires we make, the training we deliver, and the architectures we design. Precise words produce precise engineers. Vague words produce vague engineers.

◆ MEASURE 1 — DISTINGUISH CONFIGURATION FROM INFRASTRUCTURE IN VOCABULARY

Call "Configuration as Code" what is declarative configuration of virtual resources. Reserve the word "Infrastructure" for what it designates — the physical, tangible layer that carries everything else. This distinction is not a matter of style — it is a matter of operational precision. An engineer who knows that Terraform manages configuration, not infrastructure, knows there is a layer below that they cannot see and must account for.

◆ MEASURE 2 — NAME SKILLS SEPARATELY

Distinguish in job descriptions, skills frameworks and training plans what belongs to development (writing code that solves business problems), infrastructure engineering (mastering the physical layer through to the service layer) and reliability (designing and maintaining systems available in production). These three skills can coexist in one individual — some SREs excel in all three. They cannot be assumed present on the strength of a generic title.

◆ MEASURE 3 — TRAIN ON THE PHYSICAL LAYER BEFORE THE ABSTRACT LAYER

The Foundation of Iron — 26 weeks from electricity to networking — is built on this principle: you cannot master the abstraction if you do not understand what it abstracts. An engineer who learned to cable a switch before learning to configure a VPC understands what a network actually is. An engineer who only ever knew the VPC believes a network is a JSON file. Training on the physical layer is not nostalgia — it is the condition of real understanding.

Let us take back control of words before trying to take back control of our infrastructures. The two are linked. In that order.

NEMO SUPRA LEGEM EST

RAITI

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE INFRATIONAL LOOP

How DevOps manufactures the shortage that justifies its own excesses

RATIO

◆ CONTEXT OF THE STUDY

This study documents a systemic loop that has not yet been named as such: DevOps, by diluting bare-metal competences, creates a profile shortage that makes hiring salaries explode, which produces salary compression that pushes seniors to leave, which deepens the shortage, which pushes salaries back up — and so on. The direct victim is the senior engineer in post. The structural victim is digital sovereignty. The indirect beneficiary is the hyperscaler.

Amine RAITI — Infrastructure Architect & SRE

Teaching since 2006 · Electricity to Kubernetes · All audiences

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE LOOP — ANATOMY OF A SELF-CLOSING SYSTEM

SIX STEPS. ONE CYCLE. NOBODY AT THE WHEEL.

The infrational loop is not the result of malicious intent. It is the result of a series of locally rational decisions that globally produce an irrational system. Each actor acts according to their interests. The collective result is a spiral that nobody steers.

◆ THE SIX STEPS OF THE LOOP

Step 1 — DevOps dilutes bare-metal competences: documented in "The Infrational Crisis". The title "DevOps engineer" replaces "systems administrator" and "SRE" in job descriptions and training. The physical layer disappears from frameworks.

Step 2 — The bare-metal profile shortage sets in: documented in "The State and the Invisible Body". The market critically lacks engineers capable of mastering the physical layer. The shortage is structural, not cyclical.

Step 3 — Hiring salaries explode: companies needing a real bare-metal SRE — OVHcloud, Ecritel, Scaleway, but also large IT services firms — compete for a reduced pool. To attract, they outbid on entry salaries. The external market rises.

Step 4 — Salary compression is created: while the external market rises 8 to 12% per year, internal increases stagnate at 2 to 4%. Over three to five years, the new hire earns as much, sometimes more, than the senior with five years of tenure.

Step 5 — The senior leaves or complies: noticing compression, the senior engineer has three options — accept the injustice and disengage, request a pay rise (often refused or granted as a token), or leave. Most leave. Some take a deliberate voluntary gap — a few months elsewhere to "synchronise" with the market, then return elsewhere or to the same employer at market rate.

Step 6 — The shortage deepens: the senior's departure worsens the shortage. The company recruits their replacement at market rate — more expensive than what the senior would have cost after a pay rise. It has lost the tacit knowledge, it pays more to replace it, and the cycle restarts.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHEN THE NEW HIRE OVERTAKES THE SENIOR IN POST

Pay compression is a documented phenomenon in HR and labour economics. It describes the situation where the gap between experienced employees' salaries and those of new recruits narrows to the point of disappearing, or even inverting. In the tech sector in general, and in infrastructure in particular, this compression reached critical levels between 2021 and 2026.

◆ ALERT THRESHOLDS — DOCUMENTED FIGURES

Compensation management research identifies precise thresholds. Less than 10% gap between a new hire's salary and that of a 3-year employee in the same role: danger signal. Less than 5% gap: active retention risk. Negative gap — the new hire earns more than the senior: immediate flight of top performers. The cost of a key employee's departure represents 50 to 150% of their annual salary in replacement costs — recruitment, onboarding, productivity loss, tacit knowledge loss. Compression is never a long-term economy.

◆ THE MECHANICS IN BARE-METAL INFRASTRUCTURE

In France in 2026, a confirmed SRE earns on average €67,500 according to PayScale — versus €43,500 for a DevOps engineer (same source). The €24,000 gap reflects bare-metal competence scarcity. But this gap does not benefit the senior in post — it benefits the new recruit on the external market. The SRE engineer with five years of tenure in a company receiving 3% annual increases rapidly finds themselves below market rate. Their employer pays market rate to recruit their potential successor — without ever having paid this rate to retain them.

◆ THE 2021-2026 PERIOD — SYSTEMIC COMPRESSION

Intense competition for rare tech profiles has pushed entry salaries up 8 to 12% per year in the infrastructure sector, while existing employees received 2 to 4% increases. Over five years of accumulation, the gap can exceed 30 to 40%. The engineer recruited in 2021 at €55,000 with 4% annual increases is at €65,000 in 2026. Their potential market successor is offered €75,000 to €80,000. The compression is 15 to 20% — sufficient to trigger flight.

◆ NASSIHA — THIS IS NOT AN HR MISTAKE. IT IS A STRUCTURE

Pay compression is not the result of HR team carelessness. It is the structural result of the absence of a dynamic pay policy in a market under severe tension. HR teams manage increase budgets set by leadership — typically 2 to 4% — while the external market moves at 8 to 12%. Without a deliberate decision to align, compression installs itself mechanically.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE SYMBOLIC VIOLENCE OF THE DERISORY PAY RISE

The annual performance review is the moment when salary compression becomes visible to the engineer. It is also the moment when the organisation commits, often unwittingly, one of the costliest management errors possible.

◆ THE PRE-REVIEW FORM — ADMINISTRATIVE REPULSION

Before the review, there is the form. The self-assessment of 15 to 30 questions on competences, objectives, company values, strategic alignment. The infrastructure engineer who has just managed 12 production incidents, migrated two critical servers, and documented the architecture nobody else knows — this engineer spends an hour filling in a form designed to evaluate commercial profiles. The implicit signal is clear: your work is not measurable in our system. Fill in the form anyway.

◆ THE REVIEW — THE POORLY MANAGED MOMENT OF TRUTH

The review itself lasts 45 minutes. The manager is well-meaning. The feedback is positive. "You do exceptional work, you are indispensable to the team." Then comes the pay rise: €50 gross per month. €600 per year. A 1.1% increase for an engineer whose external market has grown by 10%. The dissonance between the valorising speech and the derisory pay rise is one of the most documented forms of demotivation — it says two contradictory things simultaneously: "you are valuable" and "your value does not merit real investment".

◆ THE CALCULATION THE ENGINEER MAKES IN THEIR HEAD

The engineer leaves the review and does a simple calculation. Current salary: €62,000. Market rate: €72,000 to €75,000 based on offers they see. Their increase: €700 gross annually. The compression is €10,000 to €13,000. To close it with 1.1% annual increases, they would need 12 to 15 years. Or 3 months of gap. The calculation is trivial. The decision is made.

◆ WHAT THE ORGANISATION LOSES — THE CALCULATION IT DOES NOT DO

The organisation does not do the inverse calculation. Replacement cost for this engineer: 50 to 150% of annual salary, i.e. €31,000 to €93,000. Cost of recruiting their successor at market rate: €72,000 to €75,000 salary. Tacit knowledge loss: unmeasured but documented in "The Departure of the Last One Who Knows" as irreplaceable in the short term. Investment in the pay rise that would have avoided all of this: €10,000 per year. Choosing not to revalue is a decision that costs €100,000 to save €10,000.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE VOLUNTARY GAP — THE RATIONAL STRATEGY OF THE ABSURD

LEAVING TO RETURN AT MARKET RATE

The voluntary gap is the individual and rational response to a collectively irrational system. The infrastructure engineer who has understood the mechanics of salary compression does a simple thing: they leave for a few months, and return — or go elsewhere — at market rate. It is absurd. It is effective. And it is the system itself that produces it.

◆ THE GAP MECHANICS — HOW IT WORKS

The engineer resigns. They take 2 to 4 months of gap — either to decompress, to do short freelance or interim missions, or simply to reposition on the market. They then apply — to other companies or sometimes to the same employer — with an "updated" CV and immediate availability. They are recruited at market rate, i.e. €10,000 to €15,000 more than their previous salary. The gap lasted 3 months. They recovered in one negotiation what 10 years of annual increases would not have given them.

◆ WHY THIS IS STRUCTURALLY ABSURD

The system rewards departure, not loyalty. It rewards breaking a contract, not continuity. It rewards the absence of company knowledge — the new recruit does not know the systems, past incidents, undocumented configurations — more than the presence of that knowledge. And it punishes loyalty — the engineer who stays sees their market value increase while their salary stagnates. This is not an anomaly: it is the signal the system systematically sends to its best people.

◆ WHAT THE GAP COSTS THE ECOSYSTEM

Each voluntary gap temporarily removes a rare profile from the market. If 20% of bare-metal SREs take a 3-month gap per year — a plausible figure under severe compression — 5% of the body's total capacity is unavailable at any given moment. In a body already in structural shortage, this chronic unavailability worsens the shortage, drives up entry salaries, and deepens compression for those who stayed. The loop tightens.

◆ NASSIHA — DO NOT BLAME THE ENGINEER

The voluntary gap is a rational response to an irrational signal. The engineer who takes it is not betraying their employer — they are applying the market logic their employer imposes on them. Blaming the engineer for turnover while refusing to pay market rates is blaming water for flowing downhill. The correction is structural or it is nothing.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · WHO BENEFITS FROM THIS LOOP?

THE LOOP IS NOT NEUTRAL — IT HAS BENEFICIARIES

A systemic loop that persists over time without correction always produces beneficiaries — actors whose interests are served by the status quo. Identifying these beneficiaries is not paranoia. It is structural analysis.

◆ BENEFICIARY 1 — HYPERSCALERS

A market in shortage of bare-metal competences is a market where organisations cannot build their own sovereign infrastructure — because they lack the engineers to operate it. Dependence on hyperscaler cloud is directly proportional to the shortage of local competences. Hyperscalers did not create the infrational loop — but they benefit from it structurally. Every SRE who takes a gap or recycles themselves as a "cloud architect" is one fewer engineer capable of building outside the cloud.

◆ BENEFICIARY 2 — RECRUITMENT AGENCIES

High turnover in infrastructure profiles is a windfall for specialist tech recruitment agencies. An agency placing a senior SRE receives 15 to 25% of the annual salary — i.e. €10,000 to €18,000 per placement. A market under chronic compression produces chronic turnover that produces chronic placements. This is not a conspiracy — it is an economic interest aligned with the status quo.

◆ VICTIM 1 — THE ENGINEER IN POST

The infrastructure engineer who stays in their company is the first victim of the loop. Their salary stagnates while their market progresses. Their tacit knowledge deepens while their nominal value regresses. Their engagement deteriorates — ADP documents that salary compression produces "low motivation, reduced productivity and difficulty retaining staff". They end up leaving or resigning themselves — both are a loss.

◆ VICTIM 2 — SOVEREIGN ACTORS

OVHcloud, Ecritel, Scaleway, Clever Cloud — the European sovereign actors who depend precisely on bare-metal profiles to operate their infrastructures — are victims of the loop as much as their clients. They compete for a reduced pool at inflated salaries, lose profiles to financially stronger hyperscalers, and cannot absorb the potential demand if a digital sovereignty policy were implemented tomorrow.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · BREAKING THE LOOP — SIX LEVERS

THE CORRECTION IS STRUCTURAL OR IT IS NOTHING

Breaking the infrational loop requires simultaneous interventions at multiple levels. No isolated intervention is sufficient. Together they form a correction system as coherent as the loop it seeks to interrupt.

◆ LEVER 1 — PROACTIVE ANNUAL PAY REVIEW ON MARKET BENCHMARK

Replace the internally budget-based pay review with an annual market benchmark-based review. If the market has grown 8%, the internal increase must at minimum follow — independently of the standard 3% budget. The cost of this policy is a fraction of the turnover cost it prevents. An engineer correctly valued has no reason to take the gap.

◆ LEVER 2 — RETENTION BONUS FOR RARE PROFILES

For profiles under extreme tension — bare-metal SRE, senior systems administrator, infrastructure architect — introduce an annual or biannual retention bonus, conditional on remaining with the company. This bonus is not a salary — it is an explicit recognition of competence scarcity and its replacement cost. It must be calculated on the differential between the internal salary and the external market.

◆ LEVER 3 — REMOVE THE SELF-ASSESSMENT FORM FOR TECHNICAL PROFILES

Self-assessment on generic criteria is unsuited to technical profiles whose value is operational not commercial. Replace it with a review of incidents managed, systems maintained, knowledge transmitted — metrics that correspond to the reality of the infrastructure engineer's work. The generic form is a signal of incomprehension. The operational review is a signal of respect.

◆ LEVER 4 — NAME AND MEASURE COMPRESSION

Calculate annually the ratio between in-post employees' salaries and the hiring salary for the same role. Publish this ratio internally. Set a minimum gap threshold — for example, a senior's salary must be at minimum 15% above the entry-level hiring salary for the same role. If this threshold is breached, it automatically triggers a correction budget. What is not measured is not corrected.

◆ LEVERS 5 AND 6 — TRAINING AND PUBLIC POLICY

Lever 5 — Reduce the shortage at source: the levers documented in "The State and the Invisible Body" — Foundation of Iron in BTS programmes, France Travail convention, RNCP protected title — reduce the structural shortage fuelling compression. The more qualified bare-metal engineers there are, the less entry salaries explode, the less violent the compression.

Lever 6 — Updated digital sector collective agreements: digital sector collective agreements have not been revised to account for extreme tension on infrastructure profiles. Collectively negotiated pay scales at sector level, integrating scarcity coefficients for bare-metal competences, would mechanically reduce compression by setting more realistic floors.

◆◆◆

Saving €10,000 in pay rises to lose €100,000 in replacement costs. The infrational loop is a value-destruction machine that nobody steers. It is time to name it — and break it.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

SYNTHESIS STUDY · OPÉRATION DINDON · JUNE 2026

THE GLOBAL ANATOMY OF AMPUTATION

Complete Mechanism · Solutions by Stage

Infrastructure Body & Women's Amputation

RATIO

◆ **TWO DIMENSIONS — ONE CORPUS**

This synthesis study maps two distinct and linked phenomena. **The amputation of the infrastructure body**: the body progressively loses its vocabulary, training, visibility, protection, authority and knowledge. It affects all members. **The amputation of women from the infrastructure body**: at every general stage, women experience an additional exclusion mechanism. And before any general stage, there is a stage -1 that exists only for them — the source amputation, at age 12, before they have even had the chance to enter the body.

STAGES	SOURCE STUDIES	WATERMARK
-1 → 9	**30+**	**RATIO**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

11 STAGES · TWO DIMENSIONS · ONE CHAIN

STAGE -1 — SOURCE AMPUTATION *Women only*

At age 12, in the pedagogical void, the mental image does not form. Girls do not refuse machines — nobody invited them. This stage exists only for them. It precedes all others.

↓

STAGE 0 — COGNITIVE CAPTURE *General body*

The organisation adopts hyperscaler vocabulary and metrics. Infrastructure becomes a cost, not a substrate. The loss begins without anyone noticing.

↓

STAGE 1 — VOCABULARY AMPUTATION

"DevOps" replaces precise titles. When one can no longer name, one can no longer defend. *+ Women layer: vague titles have no visible entry door. Without a clear nomenclature, no legible path.*

↓

STAGE 2 — TRAINING AMPUTATION

CPF budgets go to cloud certifications. Bare-metal disappears from training paths. Physical competence erodes. *+ Women layer: cloud training reproduces existing male informal networks.*

↓

STAGE 3 — VISIBILITY AMPUTATION

The body becomes invisible. Uptime is only seen when it fails. *+ Women layer: doubly invisible — invisible body + invisible presence.*

↓

STAGES 4-5-6 — PROTECTION · AUTHORITY · KNOWLEDGE LOST

Wrong collective agreement. Principal SRE loses final word. Tacit knowledge leaves with seniors. At each stage: additional exclusion layer for women.

↓

STAGES 7-8-9 — EXCLUSION · VOID · LOSS OF SOVEREIGNTY

The body closes in on itself. Women have no legible door. Dependency on extraterritorial hyperscalers installs structurally. CLOUD Act, lock-in, loss of mastery.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

SYNTHESIS TABLE — STAGES -1 TO 9 · DIAGNOSIS AND INTERVENTIONS

FOR EACH STAGE: MECHANISM · SIGNAL · SOURCE STUDY · INTERVENTION

STAGE	MECHANISM	WARNING SIGNAL	SOURCE STUDY	PROPOSED INTERVENTION
-1 ♀	Pedagogical void at age 12. No contact with machines. Silent bifurcation.	Tech pathways male-dominated from year 9. Zero machines in secondary school lessons.	Deamputation at the Source · The Body That Makes · Anti-Amputation Foundation	Arduino + Laser + CNC at school. Pooled teacher fund. Parents' association. Tomorrow.
0	Cognitive capture. Infrastructure becomes "a cost". Hyperscalers set the vocabulary.	Board talks "cloud strategy" with no technical infra voice. CIO from business background.	Anatomy of the Loss · No Sovereignty Without Matter	Reintroduce the physical substrate notion.
1 ♀+	"DevOps" replaces precise titles. When one cannot name, one cannot defend.	Job offers with "DevOps" only. No objective criteria. Network-based recruitment.	The Ticket and the Talent · The Infrational Crisis · The Gentle Exit	Proposed SysOps/NetOps/OpInfra nomenclature. Conversion without rupture. One job description away.
2 ♀+	CPF → cloud certifications. Bare-metal disappears. Physical competence erodes.	100% AWS/GCP training plan. No hardware or physical networking module.	The CPF Misused · The Foundation of Iron · Anti-Amputation Foundation	Neutral certifications (LFCS, CKA). Foundation of Iron. Training budget reoriented.
3 ♀+	The body is invisible. Uptime is only seen at failure. Work doesn't exist while it works.	No infra line in the annual report. No visible dedicated budget. No name.	The Economy of Silence · The Silence of the Rooms · The Uniform of the Body	Uniform of the Body — 5 levels, 5 signals, ♂/♀ options. Daily visible signal.
4 ♀+	Employer's collective agreement (banking, health, industry) ignoring infra specificities. Unregulated on-calls.	SRE in a bank covered by banking agreement. No specific night premium.	The Digital Craftsmen · The Pager and the Body	Proposed Elastic Syntec. Sectoral primacy clause. Press convention model.
5 ♀+	Principal SRE loses final word. Technical decisions made by non-technical management. Invisible debt.	Migrations decided without formalised technical opinion. Principal SRE absent from board. Post-decision incidents.	Technical Primacy · The Invisible Debt · The Infrational Loop	Technical Primacy clause in contract. Documented veto. Permanent board presence.
6 ♀+	Tacit knowledge leaves with seniors. No structured transmission. No Runbook. No mentoring.	Senior leaves → €1,200/day consultant. No runbooks. Architecture not understood in 6 months.	The Departure of the Last One · The Mentor and the Metal · The Digital Craftsmen	Formalised mentoring. 6-month supervised rotation. Mandatory machine Runbook.
7 ♀	Body closes in. Vague titles, absent nomenclature, informal networks exclude women adults.	10% women in infrastructure body. Unreadable offers. No documented progression.	The Invisible Amputation · The Gentle Exit · Maternity as Continuity	Clear nomenclature → readable offers. Objective criteria. LFCS CPF-eligible. One rewritten job offer.
8 ♀	Bifurcation at age 12. Mental image does not form. Without physical contact with machines, the path is closed.	No robotics workshop at school. Tech classes without equipment. Tech pathways seen as male.	Deamputation at the Source · The Body That Makes · Anti-Amputation Foundation	56h over 7 years. €885 budget. Pooled teacher fund + parents' association. Tomorrow.
9	Extraterritorial hyperscaler dependency. CLOUD Act, lock-in, egress fees. National sovereignty lost.	European data accessible to US DOJ. Migration impossible (egress fees). Cloud-only competences.	The Digital Ground · The Digital Iron · No Sovereignty · MariaDB Without Trembling	8 Sept 2026 ultimatum. Refurbished Grade A. Open source solutions. Progressive repatriation.

♀+ = additional exclusion layer for women at this stage.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

TOMORROW · INTERNAL DECISION · INSTITUTIONAL

◆ LEVEL 1 — TOMORROW · ZERO BUDGET · ZERO AUTHORISATION

Applicable by any teacher, engineer or manager tomorrow, without asking anyone's permission.

STAGE -1/8 (PEDAGOGICAL VOID): The Observation Day — a Principal SRE takes a group of students for a day in a datacentre. Cost: zero. Impact: often decisive.

STAGE -1/8 (THEORY): Physics teacher integrates water-electricity analogy into existing lessons. No equipment. No authorisation.

STAGE 1 (VOCABULARY): Rewrite a job description with "SysOps L3" instead of "DevOps". One job description. Today.

STAGE 6 (KNOWLEDGE): Formalise a one-page mentoring agreement between senior and junior. Two signatures. No institution.

STAGE 5 (AUTHORITY): The Principal SRE writes their first documented veto on a technical decision they refuse. One email. One signed PDF.

◆ LEVEL 2 — INTERNAL DECISION · WEEKS TO MONTHS · HR AND MANAGEMENT

STAGE 1: Adopt SysOps/NetOps/OpInfra nomenclature as internal classification grid. HR decision. One resolution.

STAGE 2: Redirect 30% of CPF budget to neutral certifications (LFCS, CKA). Annual training plan decision.

STAGE 3: Deploy the Uniform of the Body — 5 levels, 5 signals. HR + management decision. Budget: €500-700/Lead SRE.

STAGE 5: Include the Technical Primacy clause in the Principal SRE's contract. Internal legal decision.

STAGE -1/8: Open the pooled teacher fund + parents' association vote → school gift. Two meetings. No public procurement.

◆ LEVEL 3 — INSTITUTIONAL · MONTHS TO YEARS · LAW AND CONVENTION

STAGE 4: Elastic Syntec — extension of Syntec's scope to infrastructure roles regardless of employing sector. Branch agreement or extension decree.

STAGE 9: Resiliable commits, reasonable egress fees, real portability, CLOUD Act limited to its territory. European regulation.

STAGES -1/8: Official integration of the Anti-Amputation Foundation into secondary school technology curricula. Ministerial decision.

RATIO

◆ NASSIHA — DO NOT WAIT FOR LEVEL 3 TO ACTIVATE LEVEL 1

Level 1 and 2 interventions produce real effects without waiting for level 3. The Anti-Amputation Foundation does not require a ministerial decision to be deployed in one school. Technical Primacy does not require Elastic Syntec to be written into a contract. We do not wait. We act with what we have.

Amine RAITI · CC BY-NC-SA 4.0

QUICK READING GUIDE — WHICH DOCUMENT FOR WHICH PROBLEM

IF YOU OBSERVE X IN YOUR ORGANISATION · READ Y

YOU OBSERVE...	READ...
Your infrastructure engineers are all called "DevOps" and you can no longer tell who does what	The Ticket and the Talent · The Infrational Crisis · The Gentle Exit
Your training plan is 100% AWS certifications and zero hardware	The CPF Misused · The Foundation of Iron · MariaDB Without Trembling
Your infrastructure runs but nobody knows what it contributes	The Economy of Silence · The Silence of the Rooms · The Uniform of the Body
Your SRE in a bank doesn't have the same protections as in a tech firm	The Digital Craftsmen · The Pager and the Body
Management makes technical decisions without asking the Principal SRE	Technical Primacy · The Invisible Debt · The Infrational Loop
Your senior leaves and nobody knows what they knew	The Departure of the Last One Who Knows · The Mentor and the Metal · The Replacement that Reveals
There are almost no women in your infra team and you don't know why	The Invisible Amputation · The Gentle Exit · Deamputation at the Source
You are a teacher and your students never touch a physical machine	Deamputation at the Source · The Body That Makes · Anti-Amputation Foundation
You pay exorbitant egress fees and cannot migrate	The Digital Ground · No Sovereignty Without Matter · Anatomy of the Loss
You want to understand the full corpus in one document	**This study · then the Corpus Navigation Guide**

RATIO

Amine RAITI · CC BY-NC-SA 4.0

TWO AMPUTATIONS · ONE BODY TO REBUILD

The amputation of women from the infrastructure body is not a diversity problem — it is a body problem. A body amputated of half its potential population is a weakened body. It transmits less well. It defends itself less well. It innovates less well. The question is not moral — it is structural.

◆ WHAT DISTINGUISHES WOMEN'S AMPUTATION FROM GENERAL AMPUTATION

The general amputation of the infrastructure body affects all its members — it erodes vocabulary, training, visibility, protection and authority for every engineer, man or woman. Women experience this general amputation **PLUS** an additional exclusion mechanism at each stage:

At stage 1 (vague vocabulary): men can decode titles through informal networks. Women, statistically less present in these networks, carry this decoding more alone.

At stage 3 (invisibility): invisible body + invisible presence = double invisibility.

At stage 5 (authority): the female Principal SRE must doubly prove their technical legitimacy before being heard in an arbitration.

At stage 6 (knowledge): informal transmission networks (tech evenings, non-mixed Slack channels, informal mentoring) are less accessible.

And before stage 0, there is stage -1 — which exists only for women: the bifurcation at age 12, when nobody explicitly invited them to face the machine.

◆ WOMEN-SPECIFIC INTERVENTIONS — SYNTHESIS MAP

STAGE -1 (AGE 12): Anti-Amputation Foundation — the moment alone with the machine. The explicit invitation. "I want you to try this — not him in your place, you." Cost: zero. Tomorrow.

STAGE 1 (READABLE DOOR): Clear nomenclature → job offers with objective criteria → readable door without informal network. One rewritten job description away.

STAGE 3 (VISIBILITY): Feminine options for the Uniform of the Body — three choices ①②③ within the black/white/red triptych. Same authority, freedom of form.

STAGE 6 (TRANSMISSION): Feminised grades in the Digital Craftsmen (proposal): aspirant, compagnon(ne), maître(sse). Equal access as founding condition.

STAGE 7 (ADULT ENTRY): The Gentle Exit — nomenclature creates the door. Objective criteria. LFCS CPF-eligible. The senior female DevOps who wants to become Lead SysOps has a documented path.

STAGE 8 (CHILD): Deamputation at the Source + The Body That Makes. The IKEA Effect and Self-Creation Effect (Brunneder & Dholakia, 2018): making with one's hands anchors a lasting sense of competence.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

FOUNDING FORMULAS · THE EXIT PATH

◆ WHAT THE CORPUS HAS NAMED

Infrational / Infrationnelle: devaluation of technical words through DevOps vocabulary inflation.

The Infrational Loop: self-sustaining 6-step cycle producing progressive loss of bare-metal competences.

Technical Primacy: constitutive right of the Principal SRE to the final word on technical decisions.

Deamputation at the Source: intervention at school, before bifurcation happens, so the mental image forms.

The Gentle Exit: conversion without rupture from DevOps vocabulary to nomenclature, with today's tools.

The Pooled Teacher Budget: 20 teachers × €10/month = €2,400/year — a real robotics class without asking the state for a penny.

The Pedagogical IKEA Effect: making with one's hands anchors a lasting sense of competence (Norton, Mochon & Ariely, 2012).

◆ THE FORMULAS THAT SUMMARISE THE CORPUS

"The metal precedes the code."

"No sovereignty without matter."

"The one in blue will want khaki. The one in khaki will want white. This desire is the lift. The sign is the engine."

"The 12-year-old girl who blinks an LED is the bare-metal SRE of 2035."

"The Principal SRE is not infallible. But they are the only one who measures what they decide."

"The gentle exit does not say: you were wrong. It says: here is what is missing. Come."

"Without the Craftsmen, Notre-Dame would look like La Défense."

"A body without a uniform is a body without boundaries. A body without boundaries is a body without defence."

"Technical sovereignty does not begin in a datacentre. It begins in a year-7 classroom."

◆◆◆

Amputation has a mechanism. This mechanism has stages.
Every stage has an intervention. Every intervention is available.
Some from tomorrow. Without permission. Without budget. Without waiting.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · JUNE 2026

ANATOMY OF THE LOSS

Diagnosis and Reformation

From the dissolution of hard limits to a pedagogical foundation

◆ FOREWORD — A REQUEST FOR READING

This document asks to be read with reason, not emotion. It is neither an indictment of a profession nor a personal attack on anyone practising a hybrid role between development and operations. It is a structural diagnosis: an analysis of mechanisms — behavioural, economic and cognitive — that have contributed to dissolving historical technical safeguards, without passing judgment on the professional or human value of the people who practise these roles today within the framework the market has offered them.

Criticising a mechanism is not criticising the people caught within it. That is the distinction this document strives to hold from beginning to end.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

Operation Dindon

RATIO

TABLE OF CONTENTS

◆ Section 1 · The Former Ops/Dev Boundary
◆ Section 2 · DevOps as Dilution Rather Than Bridge
◆ Section 3 · Immediate Gratification and Product Design
◆ Section 4 · Certifications as Documented Capture
◆ Section 5 · The Sixth Layer of Lock-In Applied to Teams
◆ Section 6 · The 2005 Pedagogical Programme (9 Building Blocks)
◆ Section 7 · Conclusion — From Diagnosis to Operational Lever

◆ READING METHOD

Every section that puts forward a critical thesis includes a "Nassiha" callout specifying the exact limits of the claim and what is not being claimed. This discipline is deliberate: it allows a rigorous structural diagnosis to be distinguished from an emotional accusation, and invites the reader to judge this document on the soundness of its mechanisms rather than on the emotional charge an earlier, less careful formulation might have carried.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · THE HISTORICAL CONTEXT

THE FORMER OPS/DEV BOUNDARY

Before public cloud became widespread, the boundary between operations teams (system administrators, systems engineers, what we now call SRE) and development teams was clear-cut. This boundary was not arbitrary: it rested on genuine physical competence. The Ops team knew the hardware limits of the datacentre — the power available per rack, thermal capacity, cabling constraints, physical storage limits — and pushed back, with reasons, against any request that threatened those limits.

An unreasonable provisioning request from a Dev team or an outsourced web agency hit a technical red light, set by someone who genuinely understood why the request posed a risk. This refusal created real friction, often experienced as legitimate frustration by Dev teams under commercial deadlines that Ops did not always share.

> **◆ WHAT THIS BOUNDARY GUARANTEED**
>
> The Ops/Dev boundary was not a comfort of power. It guaranteed that no provisioning decision could be made without someone who understood the physical layer having a say. The cost of this guarantee was slowness. The benefit was reliability.

It would be dishonest to present this older world as perfect. The rigidities of historical Ops also produced real inefficiencies, disproportionate delivery delays, and a sometimes excessively conservative culture. Dev teams' frustration with these delays was not unfounded. It is this legitimate frustration that opened the door to the next promise: what if friction could be removed without removing reliability?

> **◆ NASSIHA — METHODOLOGICAL NOTE**
>
> Acknowledging faults on both sides of a historical boundary does not weaken the diagnosis that follows. It makes it more honest, and therefore more solid against the first objection a reader acting in good faith might raise.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

DEVOPS AS DILUTION RATHER THAN BRIDGE

The DevOps philosophy was born, originally, of a sound intention: bringing together two professions that communicated poorly, reducing time to production, and establishing shared responsibility over an application's full lifecycle. This intention is not in question here.

What followed was a structural drift more than an orchestrated plot. The spread of public cloud made it possible to create a hybrid profile — familiar with the basics of the tools Ops used (systems, middleware) and the basics of Dev's craft (agility, a few languages) — without that hybrid profile inheriting the deep physical competence that historically justified hard limits. The result is not a bridge between two complete professions. It is a dilution of both: neither the depth of the old Ops, nor the full depth of Dev, but a cross-cutting competence sufficient to operate within the cloud abstraction without ever needing to understand what lies beneath it.

> **◆ THE PRECISE MECHANISM OF THE DISSOLUTION**
>
> As long as the physical layer remained visible — cables, racks, thermal constraints — the Ops hard limit had a concrete reference point on which to operate. Once that layer became fully abstracted behind a console and an API, the hard limit lost its footing: one cannot raise a reasoned objection grounded in physical constraints to someone who, structurally, no longer sees the physical layer at all. The abstraction did not merely relocate the competence. It removed the very ground on which the refusal could be justified.

This is the precise point that distinguishes this diagnosis from mere nostalgia for the old world. The problem is not that hard limits disappeared because people became careless. The problem is that cloud abstraction made the hard limit structurally unworkable, regardless of anyone's individual goodwill.

> **◆ NASSIHA — WHAT IS NOT BEING CLAIMED HERE**
>
> This document does not claim that today's DevOps professionals are individually responsible for this dissolution, nor that they sought it. They operate within a framework offered to them by the market, the training available, and their employers. The structural responsibility described here concerns the incentive system, not the people who operate within it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE BEHAVIOURAL MECHANISM

IMMEDIATE GRATIFICATION AND PRODUCT DESIGN

A command such as an infrastructure deployment produces an immediate, visible effect: a resource appears, a server boots, a console shows a green status. This pattern — minimal action, near-instant gratification, a sense of control — has long been documented in behavioural psychology under the name of immediate or intermittent reinforcement, and has been studied for years in the design of consumer digital interfaces: notifications, game mechanics, social networks.

This document advances a precise and limited thesis: this same immediate-gratification mechanism applies to the cloud provisioning experience. Clicking or running a command and seeing a resource appear instantly produces a sense of power and control disproportionate to the actual effort involved. The more total the abstraction, the faster the gratification, and the more invisible the real cost — financial, technical, physical — becomes at the moment of the action.

> **◆ WHAT THE THESIS DOES NOT SAY**
>
> This thesis does not say that engineers are dependent in a clinical sense, nor that they suffer from a pathology. It says that cloud product design exploits a documented gratification mechanism, in the same way other sectors (video games, social media, online gambling) have been criticised for designing interfaces that exploit this same mechanism in their users. The criticism targets the product's design and the net effect it produces at an organisation's scale, not the individual psychology of whoever uses it.

The net effect of this mechanism at a company's scale is cumulative: each small individual gratification (a fast deployment, a resource appearing) reinforces the collective habit of reaching for the abstraction rather than evaluating the structural cost. This is a product-design effect that translates into organisational behaviour — not a judgment on anyone's will.

> **◆ NASSIHA — THE FORMULATION THAT WAS DISCARDED**
>
> An earlier version of this reflection directly compared the professionals concerned to people experiencing substance addiction. That comparison was removed, because it shifted the criticism from the mechanism to the person, and risked causing hurt without adding anything to the soundness of the diagnosis. The version kept here criticises the product, never the person who uses it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

CERTIFICATIONS AS DOCUMENTED CAPTURE

The trio's certification programmes (AWS Certified Solutions Architect, Microsoft Azure Administrator, Google Cloud Professional) are funded, designed, issued and maintained by the cloud providers themselves. This is not an accusation: it is a factual description of their structure. An engineer who invests several months and several hundred euros in an AWS certification has, in effect, professional capital invested in the AWS ecosystem.

This mechanism is documented across the industry under the name vendor certification lock-in. It requires no centralised malicious intent to produce a measurable net effect: a certified professional has a rational and understandable bias toward recommending, in future architecture decisions, the ecosystem in which they invested their time and money. This bias is not a moral failing. It is a logical consequence of the incentive structure.

> **◆ MARKETING THROUGH THE PROFESSION RATHER THAN THE PRODUCT**
>
> The trio did not need to advertise its products directly to technical decision-makers. It capitalised on a market dynamic — the emergence of the hybrid DevOps profession — to structure certification pathways that, while certifying a skill, simultaneously build de facto loyalty to the ecosystem. The profession becomes the vehicle for product marketing, with no classic advertising campaign required.

This thesis is distinct from a conspiracy accusation. It does not claim the trio invented DevOps with the explicit intent of creating this mechanism. It claims that, once the DevOps dynamic was established by the market, the trio structured its certification offering to extract a measurable retention benefit from it — which is rational behaviour for an economic actor, not an anomaly.

> **◆ NASSIHA — DISTINGUISHING MECHANISM FROM INTENT**
>
> Claiming that an economic actor exploited a market dynamic to its advantage is a verifiable and reasonable thesis. Claiming that it deliberately engineered that dynamic from scratch for a hidden strategic purpose is a thesis of intent that requires specific evidence — evidence this document does not possess for this particular subject, unlike the Gemplus case where institutional sources exist.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

THE SIXTH LAYER OF LOCK-IN APPLIED TO TEAMS

The six-layer lock-in model already established in this body of work places the cognitive layer as the deepest and hardest to unlock: not a contractual lock, nor a technical lock, but a lock on mental reference frameworks. A team trained and certified within the trio's ecosystem develops, by construction, a default solution framework centred on that ecosystem.

This mechanism explains an observable and frequent phenomenon: when a company's leadership, already locked into the trio, wishes to explore a return to bare metal or a local hosting provider, internal technical teams often propose, spontaneously and without deliberate intent, solutions that remain within the trio's framework — an alternative managed service, a cost-reduction option within the same ecosystem — rather than a genuine exit option.

◆ WHY THIS IS NOT SABOTAGE

A team that, through its training and professional experience, knows only one technical framework cannot, by definition, spontaneously generate solutions in a framework it has never mastered. This is not a matter of will, personal comfort, or fear for one's job. It is a structural cognitive limit: one cannot propose what one has never learned to design.

This reframing matters because it changes the nature of the appropriate response. If the problem were a matter of individual will (comfort, fear), the appropriate response would be managerial or disciplinary. If the problem is cognitive and structural — a training gap regarding an alternative framework — the appropriate response is pedagogical: broadening the available framework, not punishing the absence of a framework that was never taught.

◆ NASSIHA — THE PRACTICAL CONSEQUENCE OF THIS DISTINCTION

This precise distinction is what justifies the next section of this document. If the lock-in is cognitive rather than voluntary, the solution is not a change in individual attitude but a training programme that rebuilds the missing framework — the technical foundation that historically allowed people to understand, and therefore to propose, architectures outside public cloud.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE REFORMATION · 1/2

THE 2005 PEDAGOGICAL PROGRAMME

What follows is not an untested theoretical proposal. It is the restitution of a real vocational training programme, completed in 2005 at a vocational training centre for jobseekers, over a six-month period. This foundation, at the time, trained technicians capable of understanding infrastructure end to end, from the bit up to the operating system, before any later specialisation. It proved itself as an entry point into the profession twenty years ago, and this document's thesis is that it remains relevant today as a structural response to the cognitive lock-in described in Section 5.

1 Electrical fundamentals — power, consumption, kVA

Why this brick matters: without understanding what available power in a rack represents, it is impossible to understand why a physical hard limit historically existed. This brick restores the concrete reference point that cloud abstraction removed.

2 Number-base conversion (binary, octal, decimal, hexadecimal)

Why this brick matters: understanding the relationship between a bit, a byte and an actual electrical pulse reconnects the software abstraction to its underlying physical reality. Without this brick, the word "bit" remains an abstract concept rather than a measurable electrical fact.

3 Boolean algebra, Karnaugh maps, logic circuit design

Why this brick matters: designing a circuit from a specification (for example, controlling traffic lights) — defining variables, building the truth table, simplifying the logic, implementing it on 74XXX-series integrated circuits — teaches that all computing logic rests on real physical gates. This is the brick that makes "logic" tangible at the material level.

4 Physical automation — Arduino, programmable logic controllers

Why this brick matters: making a program interact with a real sensor or actuator restores the intuition that all code ultimately triggers a physical event, even when multiple abstraction layers obscure this link in modern cloud infrastructure.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE 2005 PEDAGOGICAL PROGRAMME — CONTINUED

5 Microcomputing — RAM, CPU, cooling, motherboard, BIOS, cabling

Why this brick matters: disassembling, identifying and assembling the physical components of a workstation or server restores the ability to diagnose a real hardware fault — a skill that quietly disappears once all infrastructure is reduced to a web console.

6 Filesystems, disk images, physical dump and restore

Why this brick matters: understanding what a disk sector is and knowing how to restore a physical image reconnects the notion of "data" to its real physical medium, rather than to the abstract, invisible promise of durability offered by a managed storage service.

7 Client/server operating systems, deployed on bare metal

Why this brick matters: installing and administering an operating system directly on physical hardware, without a cloud provider's managed hypervisor layer, restores a complete understanding of a server's lifecycle — from hardware boot to the running application service.

8 Middleware — understanding the layer between OS and application

Why this brick matters: for profiles unfamiliar with it, this brick fills a common blind spot of the modern hybrid DevOps profile, which often masters containers without ever having needed to understand what happens between the system and the application itself.

9 OSI model and network fundamentals — VLANs, segmentation, TCP/IP, routing

Why this brick matters: most network architecture decisions in public cloud today are made through drop-down menus. Understanding the OSI model through practice — simulation, real segmentation, manual routing — restores the ability to design a network rather than simply configure it in an interface.

RATIO

◆ WHAT THIS FOUNDATION PRODUCED TWENTY YEARS AGO

This programme did not train specialists in cloud abstraction — it did not yet exist in its current form. It trained technicians capable of understanding an infrastructure end to end, from the bit to the service. This end-to-end understanding is precisely what the dilution described in Section 2 has eroded. This document's proposal is not nostalgia: it is the reactivation of a foundation that has already proven its ability to produce professionals capable of setting informed hard limits, because they concretely understood what they were protecting.

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · CONCLUSION

FROM DIAGNOSIS TO OPERATIONAL LEVER

Three distinct mechanisms have been described in this document: a behavioural mechanism where product design exploits a documented immediate gratification effect, an economic mechanism where certifications produce structural loyalty without centralised malicious intent, and a cognitive mechanism where the absence of an alternative framework prevents, by construction, the spontaneous proposal of solutions outside the trio. None of the three rests on an accusation of individual ill will.

This distinction has a direct practical consequence. If the problem were one of will, the response would be disciplinary. Since the problem is structural — product design, economic incentive, cognitive gap — the appropriate response is equally structural: a training programme that rebuilds the missing framework.

> **◆ THE ROLE OF THE RE-EQUIPPED HYBRID PROFILE**
>
> A hybrid professional trained on the foundation described in Section 6 does not lose their current DevOps skills. They complement them with a framework that now includes bare metal, physical networking, and the hardware layer. This professional then becomes capable of proposing, just as spontaneously as today's solutions within the trio's ecosystem, alternative architectures with local hosting providers or in nearby server rooms — whether in France, across Europe, or elsewhere.

This document does not claim that this training programme is, on its own, a complete or sufficient solution. It claims that it is a necessary lever, documented by a real precedent — a programme that already existed and already produced results twenty years ago — and that this lever responds precisely to the cognitive mechanism identified in Section 5, rather than to a mechanism of will that was never the true cause of the problem.

> **◆ NASSIHA — FINAL METHODOLOGICAL REMINDER**
>
> This document was built while holding one constant requirement: to criticise mechanisms and incentive systems, never people. If an earlier formulation of this reflection may have appeared to target individuals, it has been withdrawn or reformulated. The reader acting in good faith is invited to judge this text on the soundness of its mechanisms, not on the emotion that an earlier, clumsier formulation might initially have provoked.

RATIO

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

Synthesis Study · Operation Dindon · June 2026

ANATOMY OF DIGITAL PERDITION

An Invulnerable Causal Cartography

From the Classroom to the Boardroom

◆ **THE CENTRAL THESIS**

The loss of digital sovereignty is not an accident, not an isolated decision, not a calculation error. It is the predictable and documentable result of a causal chain that begins in primary school and ends at the boardroom signing for AWS. This study maps seven causal nodes, four lock-in layers, three psychological locks, and five bifurcation points. It draws on the entirety of the Opération Dindon corpus — sixty-four structural studies — as evidence.

PAGES	GRAPHS	NODES	LAYERS	BIFURCATIONS
16	6	7	4	5

Amine Raiti — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Watermark RATIO + HUMAN · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

RATIO

INTRODUCTION · THE CAUSAL CHAIN — SEVEN NODES, ONE DESTINATION

Digital perdition is not an event — it is a process. It has an origin, intermediate stages, bifurcation points, and a destination. This study maps the complete causal chain across seven nodes, drawing on the Opération Dindon corpus as documentation.

NODE 1 — **Scholastic amputation + algorithmic HR sabotage** — Sin/cos without a real circle · ATS filters eliminate neutral Linux profiles · −34% Linux jobs 2019-2023. *cf. Sir, What Is This Good For (8p) · The Maths Teacher's Bag (15p) · The Newspeak (12p)*

NODE 2 — **Four entry vectors** — Free Tier · Startup Credit $100k · Free Ingress/Paid Egress · Invisible Business SaaS. *cf. The Free Tier Bait (10p) · The Free Sample (6p)*

NODE 3 — **Four lock-in layers** — Contractual (Terms §14.12) · Software (artificial complexity) · Invisible SaaS · Material (TSMC/IME). *cf. Terms Under the Microscope (8p) · The Material Black Box (4p) · The Taiwan Bottleneck (6p)*

NODE 4 — **Geopolitical rationing** — TSMC ~90% advanced chips · blockade scenario → cloud queue · Refurbished = autonomy. *cf. The Taiwan Bottleneck (6p) · The Strategic Refurbished (4p)*

NODE 5 — **Three psychological locks** — Sunk Cost Fallacy (Kahneman 2011) · Cloud-Washing 3 layers · Consulting firm authority bias (AWS Premier Partner). *cf. The Cloud-Washing (7p) · The Free Sample S4*

NODE 6 — **Institutional perdition** — DINUM · CPF x3 · Free Tier local authorities · abandoned pedagogical infrastructure budget. *cf. The Cloud-Washing S3 · The CPF Misused (6p) · The Free Tier Bait S6*

NODE 7 — **Five bifurcation points** — At each stage, the trajectory toward perdition could have been interrupted. Cost increases from bifurcation 1 (near-zero) to bifurcation 5 (total rewrite). *→ Section 6 · Graph 6*

◆ **GRAPH 1 (NEXT PAGE) MAPS THESE SEVEN NODES AS AN ORIENTED CAUSAL TREE**

Read left to right. Colours: blue=education · red=economic · grey=material · violet=HR · orange=applicative SaaS.

RATIO

CC BY-NC-SA 4.0

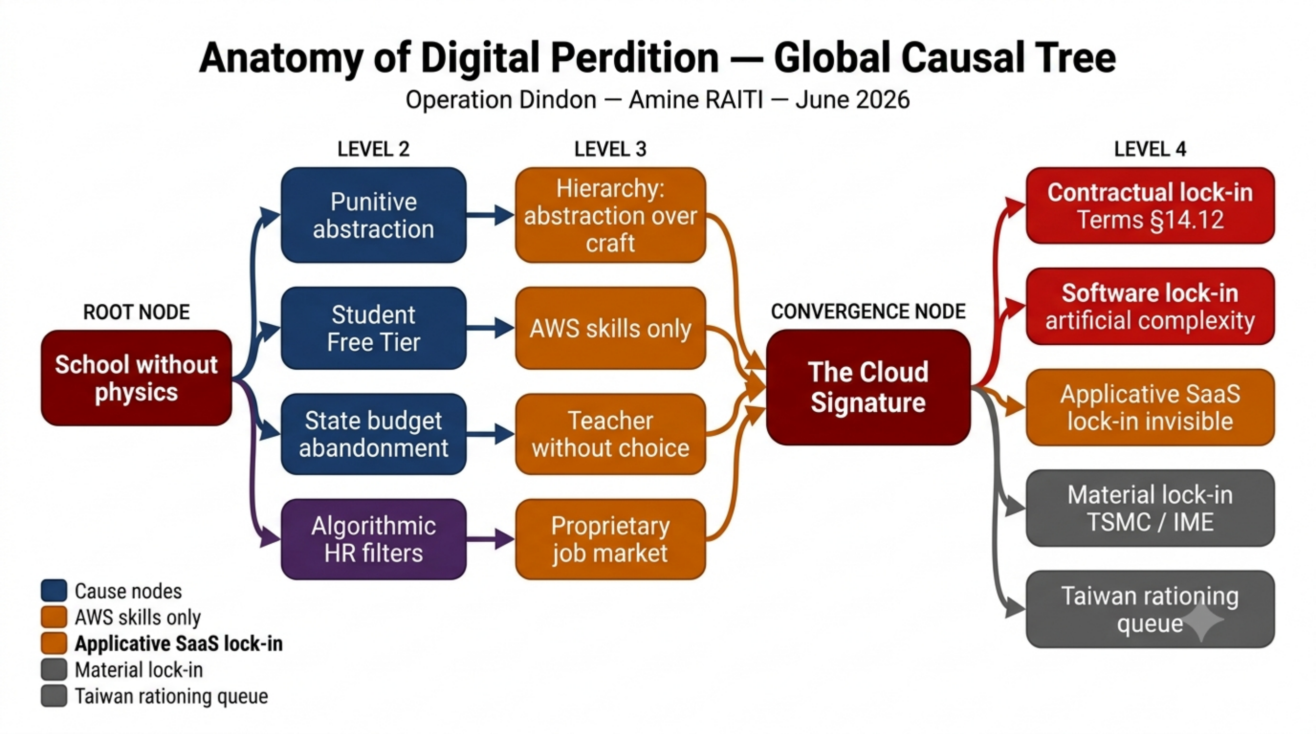

GRAPH I · Oriented DAG left→right · Colour-coded nodes by category · OPERATION DINDON · AMINE RAITI

CC BY-NC-SA 4.0

I

SECTION I · THE AMPUTATION AT THE SOURCE — SCHOLASTIC AND ALGORITHMIC

SCHOOL WITHOUT PHYSICS + ALGORITHMIC HR SABOTAGE: TWO MECHANISMS THAT REINFORCE EACH OTHER

◆ A · SCHOLASTIC AMPUTATION — CF. SIR, WHAT IS THIS GOOD FOR (8P) · THE MATHS TEACHER'S BAG (15P)

School teaches sine and cosine without ever showing a real circle. It teaches networks without ever opening a patch panel. It teaches databases without ever configuring a physical server. This rupture between symbol and matter creates a developer who knows how to deploy a Lambda function without knowing what runs underneath. Free Tier reinforces this amputation: it gives access to cloud abstractions without ever requiring understanding of fundamentals. The State's budgetary abandonment eliminated lab rooms with real servers. The teacher does not choose between AWS and a sovereign alternative — they choose between AWS and nothing. AWS Educate does not exist out of philanthropy. It exists because Amazon identified this void as an acquisition opportunity.

◆ B · ALGORITHMIC HR SABOTAGE — BLIND SPOT INTEGRATED · CF. THE NEWSPEAK (12P)

Even if a CIO decides to recruit neutral skills, recruitment platforms (LinkedIn, Indeed, internal ATS) filter CVs on hyperscaler-certified keywords: "AWS Cloud Architect", "Azure Certified Developer", "GCP Professional". A "Senior Linux Administrator" or "Bare-Metal Engineer" profile is eliminated algorithmically before the interview. Lock-in is maintained by an invisible bureaucratic barrier — not by a deliberate human decision. Documented figure: −34% Linux administrator positions between 2019 and 2023 on French job platforms.

◆ THE CLOSED LOOP — GENERATION N EFFECTS BECOME GENERATION N+1 CAUSES

The student trained on AWS becomes the senior developer who recommends AWS. They become the architect who configures ATS filters on AWS keywords. They become the CIO who recruits only AWS-certified profiles. Graph 2 (heatmap, next page) maps the causal intensity from each scholastic cause to each professional effect, with the diagonal showing generational self-reinforcement.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

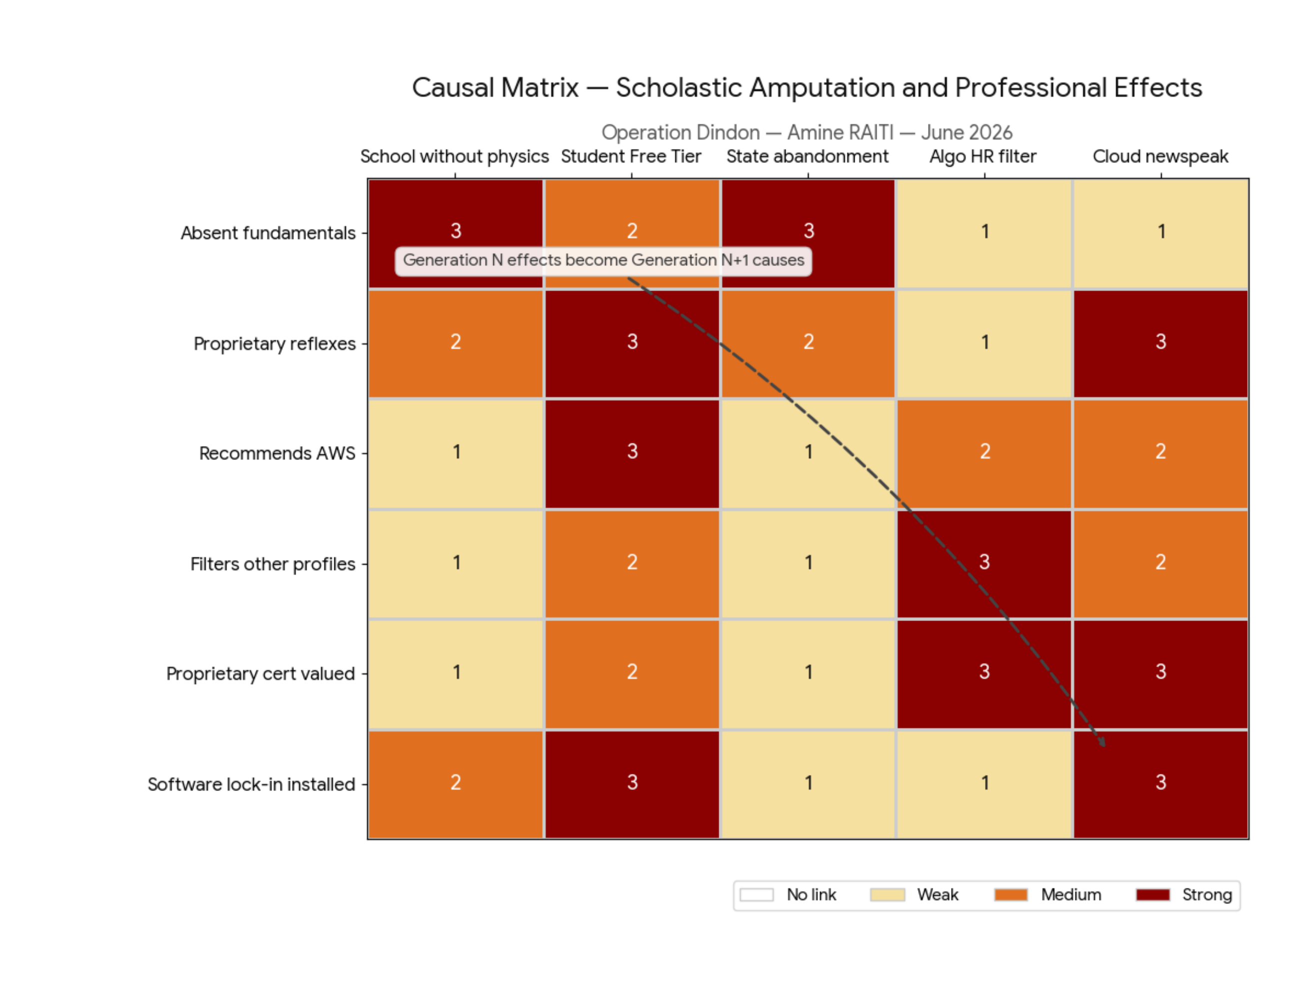
Causal Matrix — Scholastic Amputation and Professional Effects
Operation Dindon — Amine RAITI — June 2026
School without physics
Student Free Tier
State abandonment
Algo HR filter
Cloud newspeak
Absent fundamentals
Proprietary reflexes
Recommends AWS
Filters other profiles
Proprietary cert valued
Software lock-in installed
3
2
3
1
1
2
3
2
1
3
1
3
1
2
2
1
2
1
3
2
1
2
1
3
3
2
3
1
1
3
Generation N effects become Generation N+1 causes
No link
Weak
Medium
Strong

CC BY-NC-SA 4.0

2

SECTION 2 · THE FOUR ENTRY VECTORS — FREE AT FIRST · COSTLY TO LEAVE

FOUR DOORS · ONE LOGIC: FREE WHEN DEPENDENCY INSTALLS · CHARGED WHEN YOU WANT TO LEAVE

◆ DOOR 1 · STUDENT FREE TIER — CF. The Free Tier Bait (10P)

AWS Educate · Azure for Students. 750h EC2 · 5 GB S3 · 1M Lambda requests. Invisible threshold — no workflow alert. Exceeded when attachment is maximal and migration most costly.

◆ DOOR 2 · STARTUP CREDIT — CF. The Free Sample (6P)

AWS Activate $100,000 · Google for Startups $200,000. Architecture without financial constraint → TCO crisis at exhaustion. VC — startup — hyperscaler triangle: the first two fund the third.

◆ DOOR 3 · DATA GRAVITY TRAP — CF. The Free Tier Bait S5

Ingress: free · Egress: $0.09/GB AWS. 100 TB = $9,000 exit ransom. 1 PB = $90,000. Entry is free because it creates dependency. Exit is charged because it reveals its cost.

◆ DOOR 4 · INVISIBLE BUSINESS SAAS — Blind spot integrated

HR buys Workday · Marketing buys Salesforce · Sales buys HubSpot — without IT. 73% of organisations have SaaS purchased outside IT. Repatriating servers does not resolve this applicative layer.

◆ THE 4-STEP SEQUENCE — COMMON TO ALL 4 DOORS — CF. The Free Sample S2

1. Free entry sample — no financial friction. 2. Functional dependency — integration into workflows and architectures. 3. Tipping threshold — the free stops. 4. Exit friction — the captive ecosystem makes return to independence costly. Graph 3 (funnel, next page) maps these four converging doors with bifurcation points at each stage.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

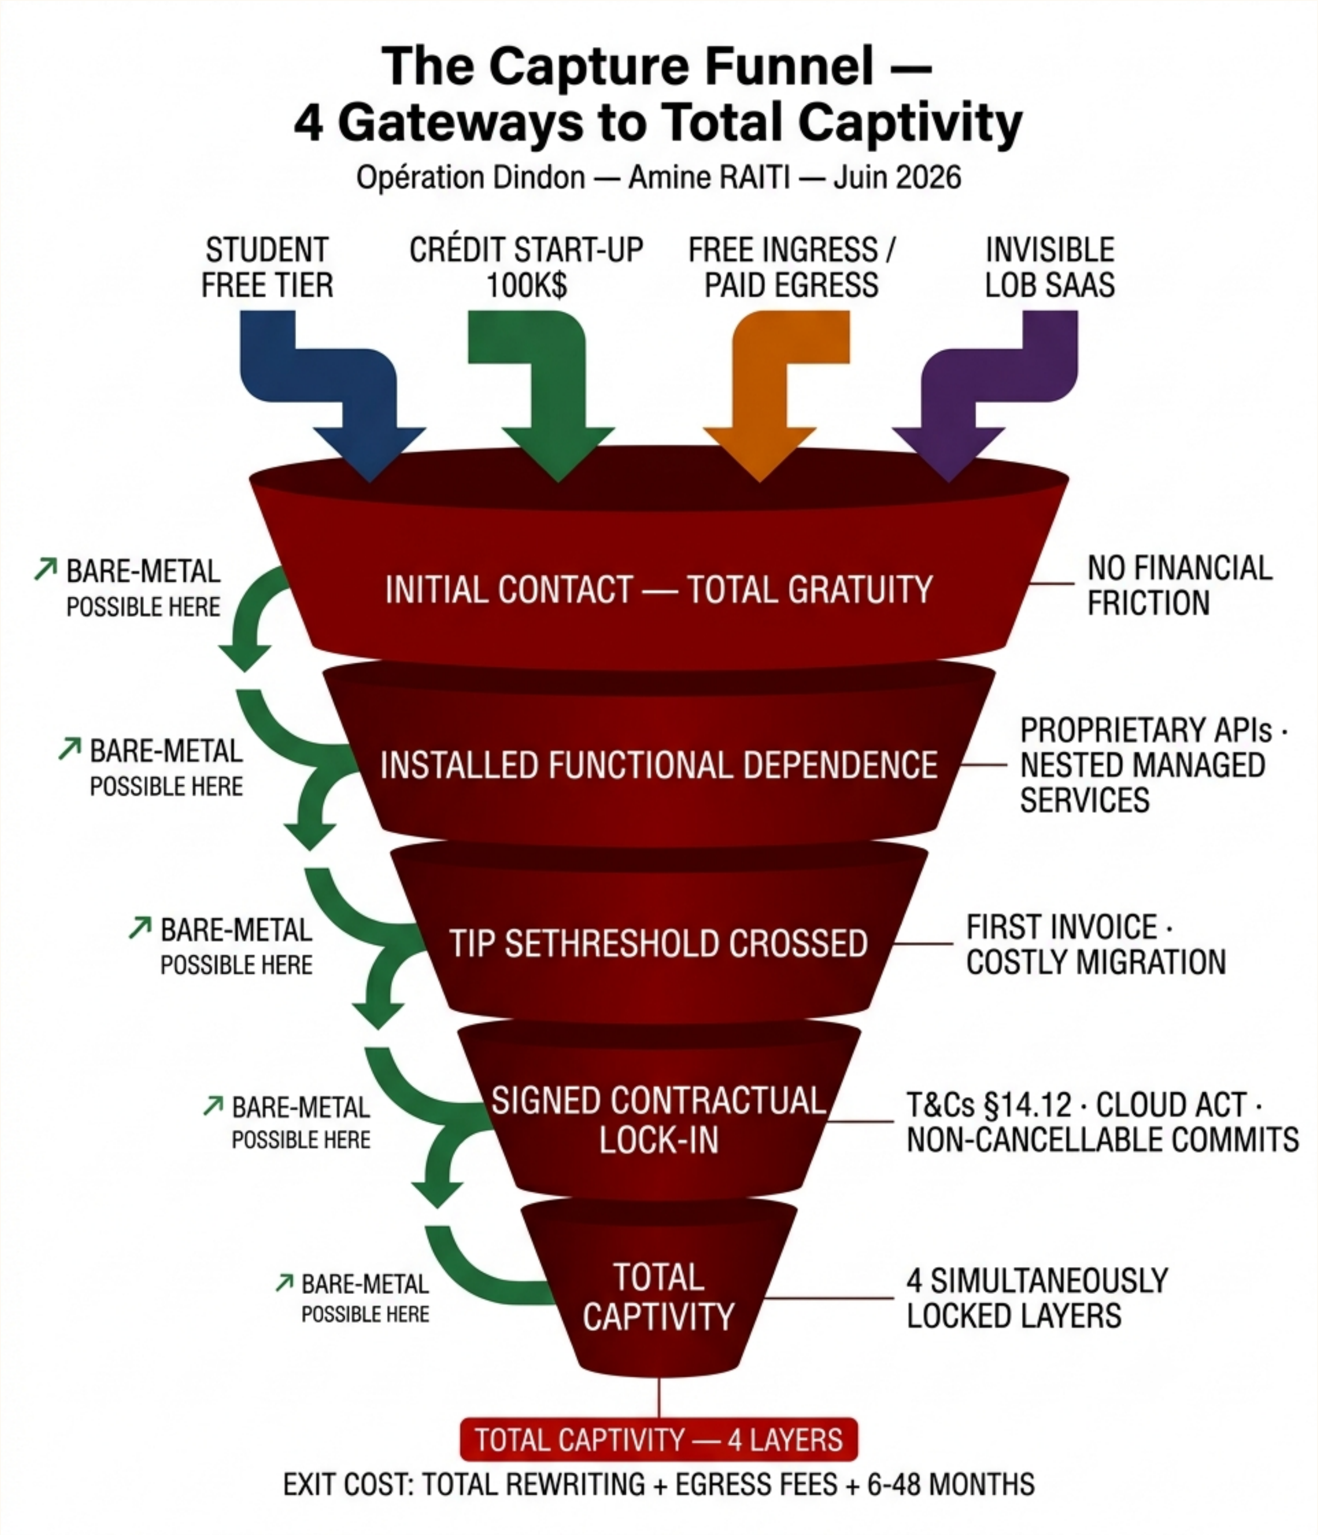

GRAPH 3 · Vertical funnel chart · 4 entries · 5 stages · lateral bifurcations · OPERATION DINDON · AMINE RAITI

CC BY-NC-SA 4.0

CONTRACTUAL · SOFTWARE · APPLICATIVE SAAS · MATERIAL — INCREASING DEPTH AND EXIT COST

◆ LAYER 1 — CONTRACTUAL — cf. Terms Under the Microscope (8p) · Exit cost: MEDIUM · 6-18 months

Terms §14.12: California jurisdiction · CLOUD Act applicable · non-cancellable commits · egress fees. Sovereignty Washing (cf. The Cloud-Washing S1) sells physical resolution as if it also resolved the legal dimension. "AWS European Sovereign Cloud" = servers in Europe (physical ✓) but Amazon.com Inc. remains subject to CLOUD Act (legal ✗). The four demands of the 8 September 2026 ultimatum target exactly this layer.

◆ LAYER 2 — SOFTWARE — cf. The Newspeak (12p) · The Free Sample S5 · 12-36 months

Lambda + DynamoDB + API Gateway + SQS + Cognito + CloudFront — each link does not exist outside AWS in this exact form. The chain is proprietary. Migrating = total rewrite. Cloud-Washing calls this "cloud-native." Architectural question: "*Can this component run outside* AWS *without a rewrite?*" If no — artificial complexity, not modernity.

◆ LAYER 3 — APPLICATIVE SAAS — Blind spot integrated · 18-48 months

73% of organisations have SaaS purchased outside IT. Salesforce = Sales business logic · Workday = HR history · M365 = the entire organisation's documents. Repatriating servers resolves the infra layer — not the applicative layer. Two sovereignties to reconquer simultaneously, the second involving business directorates without IT constraints.

◆ LAYER 4 — MATERIAL — cf. The Taiwan Bottleneck (6p) · The Material Black Box (4p) · Horizon 10-15 years

TSMC fabricates ~90% of advanced chips worldwide in the Taiwan Strait. Intel ME / AMD PSP: non-auditable sub-processors since 2008/2013, subject to US law, active with server off. Taiwan blockade: hyperscalers will ration in order — internal needs → Enterprise Agreement → Pay-as-you-go. Without owned hardware: queue. With Grade A Refurbished already fabricated (cf. The Strategic Refurbished, 4p): complete autonomy. Chips exist, independent of future TSMC production. Potential blockade impact: $600-1,200 billion (PIIE 2023).

RATIO

Amine RAITI · CC BY-NC-SA 4.0

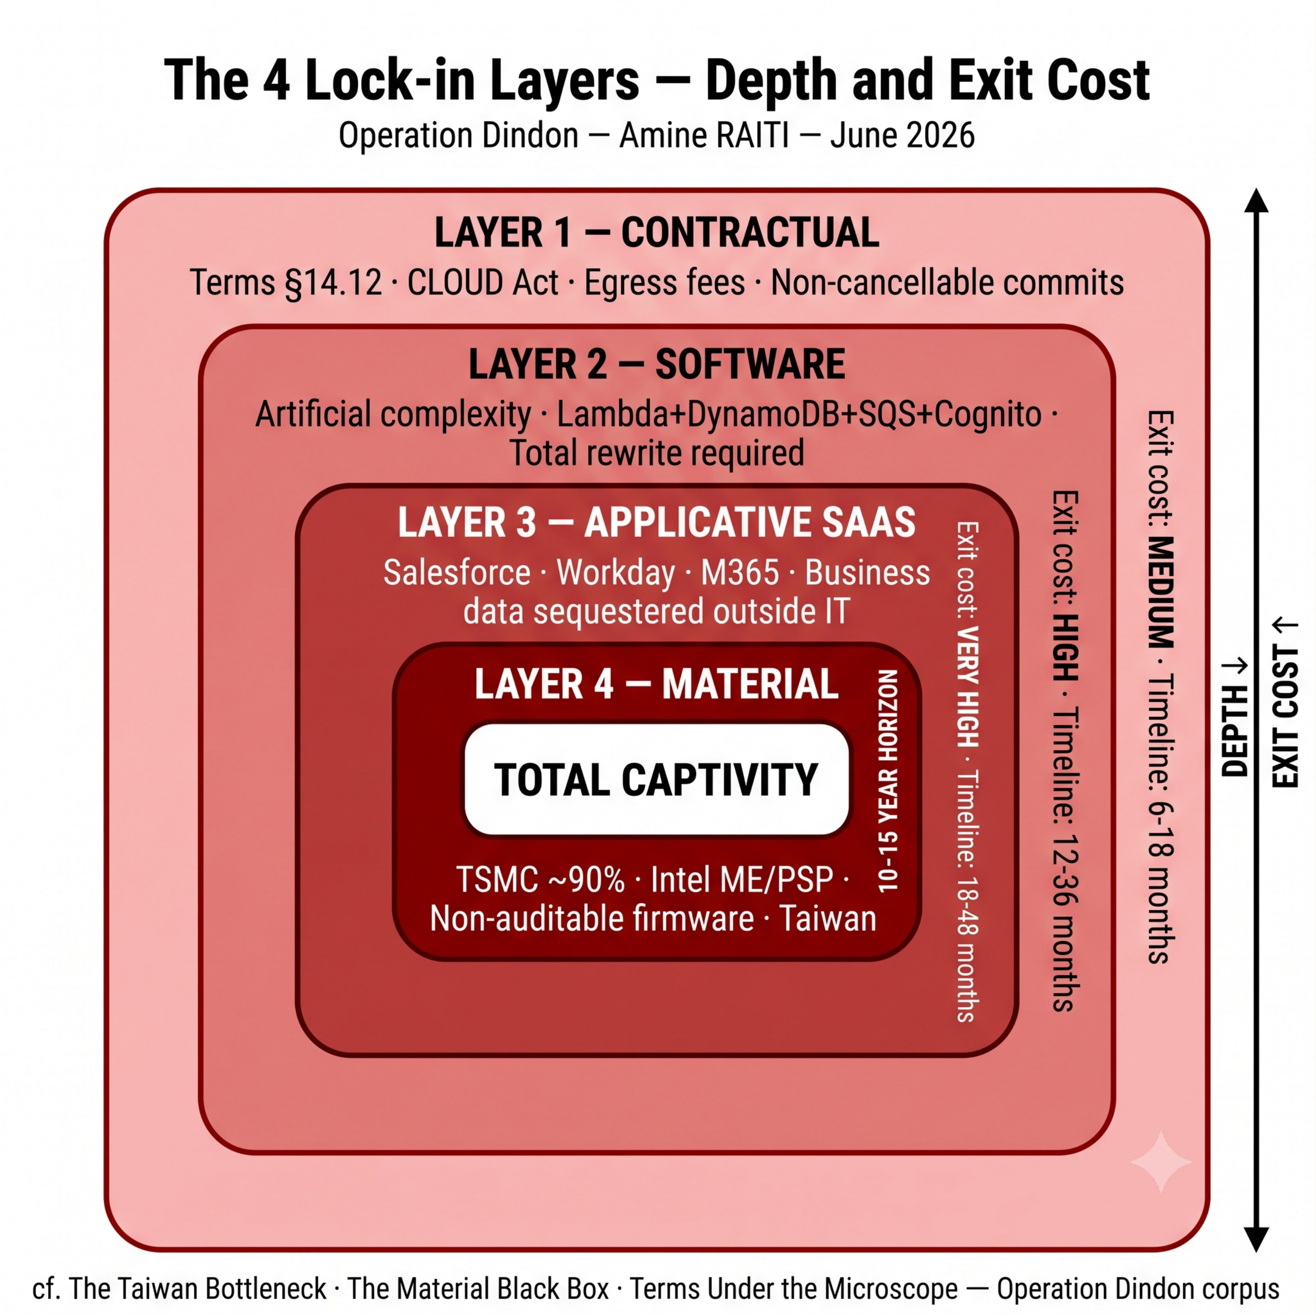
The 4 Lock-in Layers — Depth and Exit Cost
Operation Dindon — Amine RAITI — June 2026
LAYER 1 — CONTRACTUAL
Terms §14.12 · CLOUD Act · Egress fees · Non-cancellable commits
LAYER 2 — SOFTWARE
Artificial complexity · Lambda+DynamoDB+SQS+Cognito · Total rewrite required
LAYER 3 — APPLICATIVE SAAS
Salesforce · Workday · M365 · Business data sequestered outside IT
LAYER 4 — MATERIAL
TOTAL CAPTIVITY
TSMC ~90% · Intel ME/PSP · Non-auditable firmware · Taiwan
10-15 YEAR HORIZON
Exit cost: VERY HIGH · Timeline: 18-48 months
Exit cost: HIGH · Timeline: 12-36 months
Exit cost: MEDIUM · Timeline: 6-18 months
DEPTH →
EXIT COST ↑
cf. The Taiwan Bottleneck · The Material Black Box · Terms Under the Microscope — Operation Dindon corpus

CC BY-NC-SA 4.0

4

SECTION 4 · THE THREE PSYCHOLOGICAL LOCKS — WHAT DATA ALONE CANNOT UNLOCK

SUNK COST FALLACY · CLOUD-WASHING · CONSULTING FIRM AUTHORITY BIAS

◆ LOCK 1 · SUNK COST FALLACY — KAHNEMAN 2011 · THALER 1980 · CF. THE FREE SAMPLE S4

The CIO who evangelised cloud for 5 years on LinkedIn, received a promotion for the digital transformation, convinced their CEO — cannot retreat without publicly admitting they were wrong. The fear of losing face before peers is stronger than economic rationality. The exit narrative: "*The 2019 decision was correct with 2019 information. 2026 information is different — CLOUD Act clarified, egress fees documented, TSMC geopolitical risk identified. Revising the position is good management.*" Not communication — the truth.

◆ LOCK 2 · CLOUD-WASHING — CF. THE CLOUD-WASHING (7P) · MU'ALLAQA IV · AMINE RAITI

Three layers neutralising the three internal resistances: **SOVEREIGNTY WASHING** reassures legal (servers Europe ✓ / California jurisdiction ✗) · **GREEN WASHING** reassures sustainability (RECs ≠ green energy · Virginia 30-35% fossil EIA 2023) · **FINANCE WASHING** reassures CFO (public money → hyperscalers → anti-regulation lobbying · CPF x3). The engineer who says "let us calculate the real TCO" stands alone against three arguments, three directorates, three boxes already ticked. Cloud-Washing is not an isolated lie — it is an architecture of conviction.

◆ LOCK 3 · CONSULTING FIRM AUTHORITY BIAS — BLIND SPOT INTEGRATED

McKinsey, Accenture, Capgemini are "AWS Premier Partner", "Google Cloud Partner". Their business model depends on billing complex migration person-days. The simplicity of local Bare-Metal generates no consulting revenue. A Gartner report recommending AWS is funded, directly or indirectly, by the hyperscaler ecosystem. The CIO receives an "objective" report confirming what Cloud-Washing already told them. Two external validation systems — and both have structural conflicts of interest.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

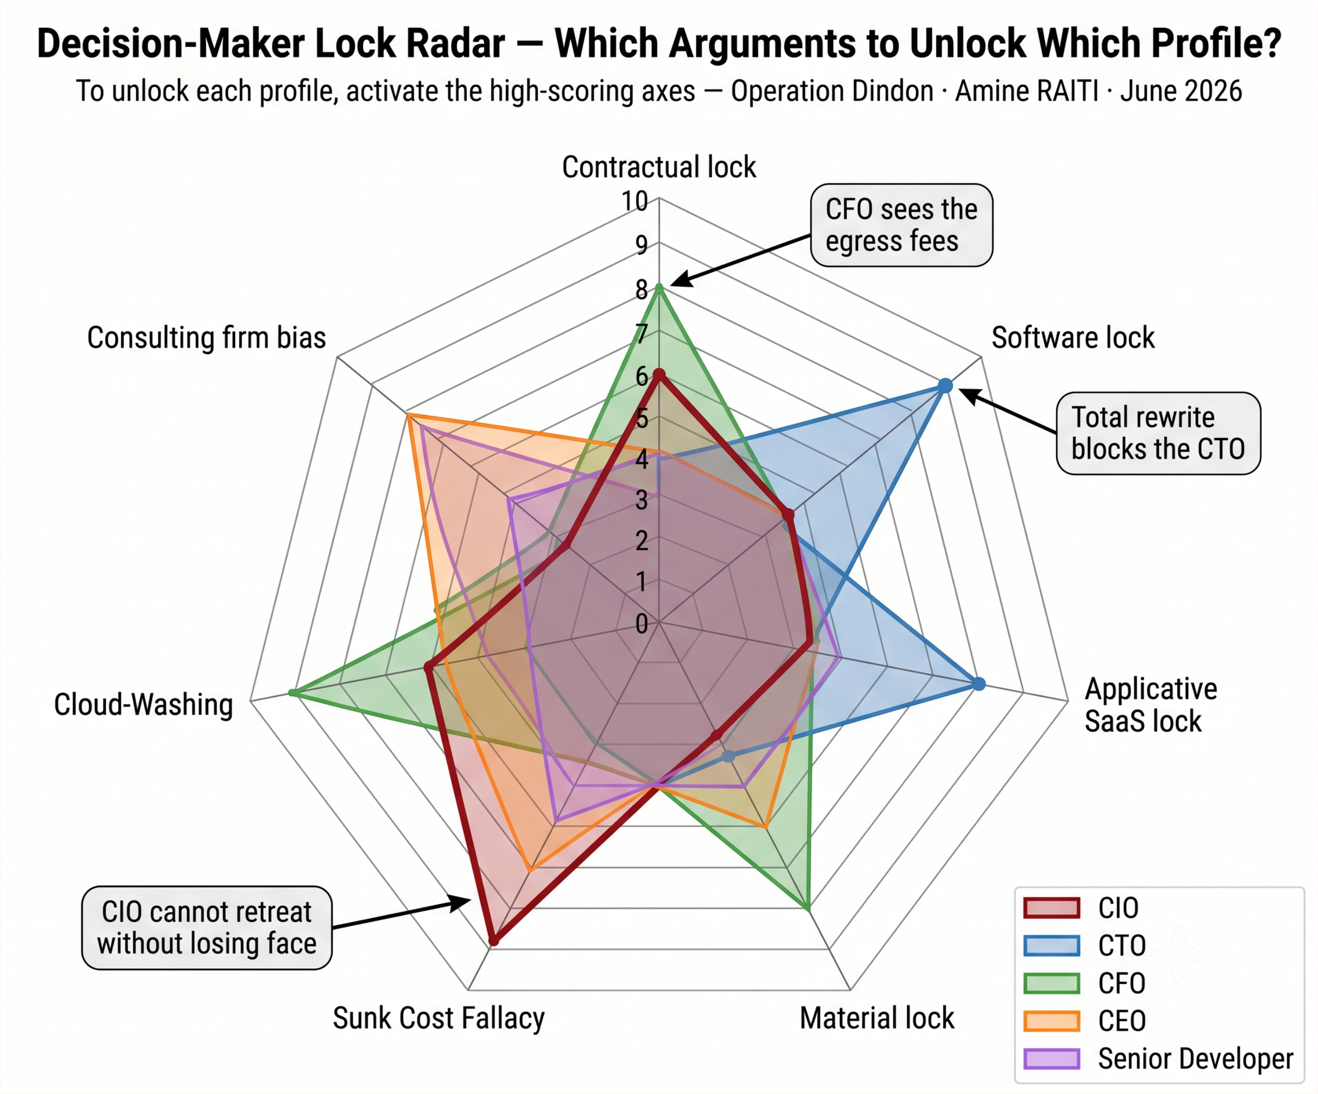
Decision-Maker Lock Radar — Which Arguments to Unlock Which Profile?
To unlock each profile, activate the high-scoring axes — Operation Dindon · Amine RAITI · June 2026
Contractual lock
Software lock
Applicative SaaS lock
Material lock
Sunk Cost Fallacy
Cloud-Washing
Consulting firm bias
0
1
2
3
4
5
6
7
8
9
10
CFO sees the egress fees
Total rewrite blocks the CTO
CIO cannot retreat without losing face
CIO
CTO
CFO
CEO
Senior Developer

CC BY-NC-SA 4.0

5

SECTION 5 · INSTITUTIONAL PERDITION — FOUR LEVELS OF PUBLIC MONEY ABSORPTION

THE TAX THAT FUNDS THE DEPENDENCY IT CLAIMS TO COMBAT

◆ LEVEL 1 · PUBLIC PROCUREMENT — CF. THE CLOUD-WASHING S3

DINUM selects AWS, Azure and Google for French public administrations. Public hospitals migrate to Azure. Universities use Google Workspace. Ministries deploy on AWS. The EU uses US hyperscaler services while funding GAIA-X. Paradox: public money funds hyperscalers through procurement → same hyperscalers use this money for lobbying on GDPR, Data Act, AI Act → tax funds resistance to the regulations that would protect the taxpayer.

◆ LEVEL 2 · PUBLIC TRAINING FUNDS — CF. THE CPF MISUSED (6P)

Cloud certification spending tripled between 2020 and 2023. Public professional training money — salary contributions — massively funds AWS, GCP and Azure certifications. These certifications reinforce dependency by creating a labour market that values proprietary skills above neutral skills. The public training system funds the cognitive capture documented in The Newspeak.

◆ LEVEL 3 · FREE TIER IN LOCAL AUTHORITIES — CF. THE FREE TIER BAIT S3

Local authority "tests" on AWS Free Tier — "free to start." Six months later: paying customer of a US hyperscaler for a public service funded by local tax. No procurement process — project below public contract thresholds. The taxpayer will pay indefinitely or invest in a costly migration.

◆ LEVEL 4 · ABANDONED PEDAGOGICAL INFRASTRUCTURE — CF. THE FREE TIER BAIT S6

The State eliminated physical lab budgets. AWS Educate filled this void — not out of philanthropy but because Amazon identified an acquisition opportunity. **THE CLOSED LOOP:** tax → eliminates TP budget → hyperscaler fills the void → student trained on AWS → CPF funds AWS certifications → CPF is funded by tax → loop. Every euro of cloud CPF reinforces the causal chain documented in this study.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

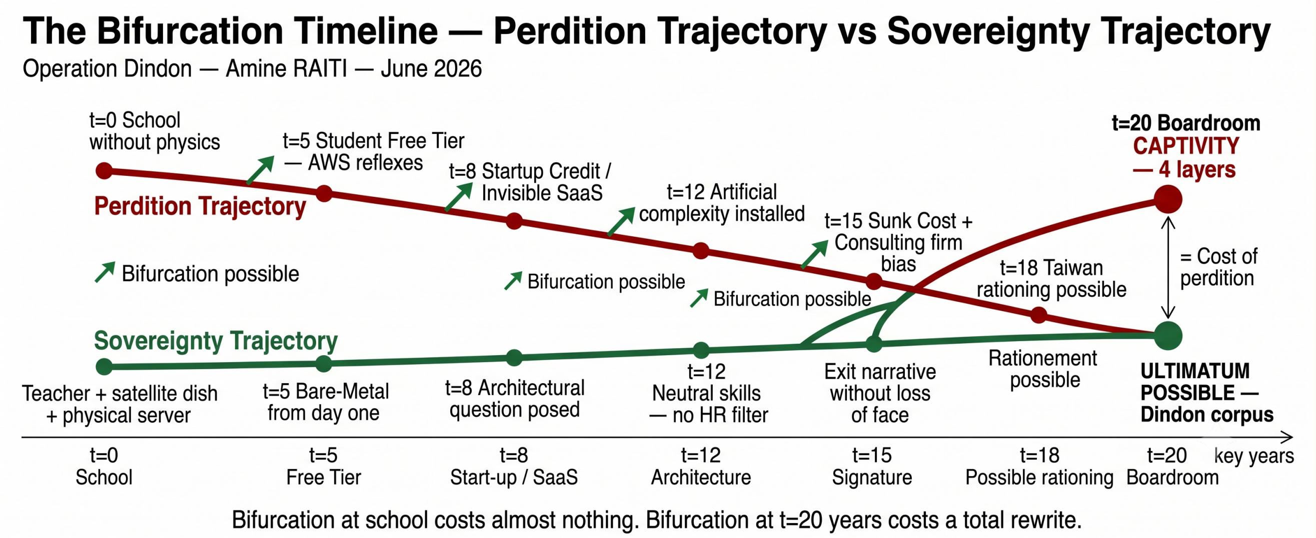

GRAPH 6 · Horizontal timeline · 2 trajectories · 5 bifurcation points · OPERATION DINDON · AMINE RAITI

CC BY-NC-SA 4.0

6

SECTION 6 · THE FIVE BIFURCATION POINTS — WHERE THE CHAIN COULD HAVE BEEN INTERRUPTED

THE LATER THE INTERVENTION · THE HIGHER THE COST · THE LAW OF LATE BIFURCATION

#	MOMENT	ACTION	EFFECT	COST
1	**t=0 · School**	Satellite dish + physical server in university · The Maths Teacher's Bag	Fundamentals acquired · Chain interrupted at source · HR filters neutralised	**Near-zero**
2	**t=5 yrs · Free Tier**	Refuse Free Tier · Bare-Metal from day one · Architectural question	No functional dependency · TCO visible immediately · No invisible threshold	**Immediate cost · Better TCO over 5 years**
3	**t=8 yrs · Architecture**	'Can this component run outside AWS without a rewrite?' · Refuse artificial complexity	No software lock-in · Migration possible without total rewrite	**Zero — architectural decision**
4	**t=15 yrs · Signature**	Read Terms §14.12 · Calculate real TCO · Exit narrative without loss of face	Conscious decision · Ultimatum possible · Dindon corpus as documentation	**Time + legal expertise**
5	**t=20 yrs · Perdition**	Accompanied Exodus · Sovereign Interface · Gentle BCP · Progressive repatriation	Return to independence — difficult but possible · 4 layers to unlock	**6-48 months · Partial rewrite · High cost**

◆ THE LAW OF LATE BIFURCATION

Bifurcation 1 (school): near-zero budget. Bifurcation 3 (architecture): zero. Bifurcation 5 (perdition): 6-48 months + rewrite. Each uninterrupted causal node exponentially increases the cost of the next. The chain is predictable, documented, and avoidable — provided intervention comes early.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

15 MECHANISMS · 64 CORPUS STUDIES · THE MAP AND ITS EVIDENCE

MECHANISM	CORPUS STUDY	KEY EVIDENCE	BIFURCATION
Scholastic amputation · sin/cos without circle	*Sir, What Is This Good For (8p) · The Maths Teacher's Bag (15p)*	Symbol/matter rupture · teacher without choice	Physical server in university
Algorithmic HR sabotage	*The Newspeak (12p)*	−34% Linux jobs 2019-2023	Recruit on neutral skills
Student Free Tier · AWS Educate	*The Free Tier Bait (10p)*	750h EC2 · invisible threshold · bill shock	Bare-Metal from day one
Startup Credit · TCO Crisis	*The Free Sample (6p)*	$100k AWS Activate → architecture without constraint	Size to real revenues
Data gravity trap Ingress/Egress	*The Free Tier Bait S5 · Terms Under the Microscope (8p)*	100 TB ingress free · egress $9,000	Data on own hardware from start
Invisible Business SaaS (shadow IT)	*Blind spot — original contribution*	73% SaaS outside IT · Salesforce · Workday	IT SaaS policy before purchase
Contractual lock-in Terms §14.12	*Terms Under the Microscope (8p)*	California jurisdiction · CLOUD Act · non-cancellable commits	Read Terms · ultimatum 4 demands
Software artificial complexity	*The Free Sample S5 · The Newspeak*	Lambda+DynamoDB+SQS+Cognito = total rewrite	'Runs outside AWS without rewrite?'
Material lock-in · TSMC · IME/PSP	*The Taiwan Bottleneck (6p) · The Material Black Box (4p)*	TSMC ~90% advanced chips · IME non-auditable	Grade A Refurbished already fabricated
Geopolitical rationing Taiwan	*The Taiwan Bottleneck (6p) · The Strategic Refurbished (4p)*	Blockade → cloud queue vs hardware autonomy	Owned hardware = outside queue
CIO Sunk Cost Fallacy	*The Free Sample S4 · Kahneman 2011*	LinkedIn + promotion = publicly invested ego	Exit narrative without loss of face
Cloud-Washing 3 layers	*The Cloud-Washing (7p) · MU'ALLAQA IV*	Sovereignty ✓ physical ✗ legal · Virginia 30-35% fossil	Map 3 layers before decision
Consulting firm authority bias	*Blind spot — original contribution*	AWS Premier Partner = structural conflict of interest	Independent TCO audit mandatory
Institutional perdition · DINUM · CPF	*The Cloud-Washing S3 · The CPF Misused (6p)*	CPF x3 cloud certifications 2020-2023	Public procurement: sovereign alternatives
Free Tier local authorities without procurement	*The Free Tier Bait S3*	Authority paying customer without deliberate decision	Procurement thresholds from Free Tier

This table is the entry point of the Opération Dindon audit library. Each row points to the corresponding 4-to-15-page evidence study. Two mechanisms (Invisible Business SaaS · Consulting firm bias) are original contributions of this synthesis study.

RATIO

CC BY-NC-SA 4.0

CLOSING · THIS IS NOT A CONCLUSION — IT IS A FINDING

The chain is documented. The mechanisms are named. The locks — technical, financial, psychological — are mapped. The five bifurcation points are identified with their exact costs. What remains: the decision.

Digital perdition is not inevitable — it is the result of a chain of absences. The absence of a satellite dish at school. The absence of a physical server in the lab. The absence of reading the Terms. The absence of the architectural question. The absence of a procurement process before "testing" Free Tier. Each absence is a choice — often unconscious, always predictable.

The Opération Dindon corpus is the documentation. The ultimatum of 8 September 2026 is the demand. Four requirements: cancellable commits · elimination of egress fees · guaranteed data portability · clarification of the CLOUD Act position. The response from AWS, GCP and Azure is the unknown variable.

HUMAN

The Bare-Metal laughs at your labels of card.

For the Iron lies not, 'tis our staff and our guard.

Amine RAITI · MU'ALLAQA IX · 2026

NEMO SUPRA LEGEM EST

Operation Dindon · Amine RAITI · Infrastructure Architect & SRE · Teaching since 2006

CC BY-NC-SA 4.0 · June 2026 · 64 structural studies · Ultimatum: 8 September 2026

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026

THE DISSOLUTION OF THE BOUNDARY

How the DevOps Merge Accelerated Perdition

Ego · Temporalities · IaC · Serverless · FinOps · Sovereign IDP

◆ THE THESIS

The Operation Dindon corpus documented all external locks — Terms, egress fees, TSMC, Cloud-Washing. This study documents the internal lock that precedes all others: the dissolution of the boundary between development and infrastructure. Not a nostalgic defence of territory. A finding that two professions with fundamentally different temporalities, risks and natures were merged under the same umbrella — and that this merger is one of the intimate drivers of cloud captivity. The first act of sovereignty is not technical. It is recognising that we were wrong.

HUMAN

PAGES	BLIND SPOTS	WATERMARK	GRAPHS
8	5	HUMAN DOMINANT	0

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · *Operation Dindon* · *June 2026*

CC BY-NC-SA 4.0

I

SECTION I · EGO AS THE FIRST LOCK — DEFEATING FEAR BEFORE DEFEATING LOCK-IN

THE FIRST ACT OF SOVEREIGNTY IS NOT TECHNICAL — IT IS RECOGNISING THAT WE WERE WRONG

◆ THE DECISION-MAKER'S SUNK COST FALLACY — cf. The Free Sample S4 · Kahneman 2011 · Thaler 1980

A CIO who spent five years evangelising the DevOps transformation of their organisation — merging teams, eliminating "silos," celebrating deployment velocity — cannot retreat without publicly admitting they were wrong. Not cowardice. The Sunk Cost Fallacy documented by Kahneman (2011) and Thaler (1980): the tendency to continue a commitment not because it is rational, but because one has already invested too much to turn back.

The investment is double. Financial first — migrations, training, certifications. But above all public — LinkedIn posts, conferences, internal use cases, the professional reputation built on this transformation. Retreating means erasing five years of professional posture before peers and the CEO. The fear of losing face is stronger than economic rationality. This is the lock that blocks the cloud exit before a rack has even been opened.

◆ THE EXIT NARRATIVE WITHOUT LOSS OF FACE

The decision to merge infra and dev teams in 2019 was correct with 2019 information. The DevOps movement had real arguments — breaking delivery silos, reducing time to market, making teams responsible for quality in production. These arguments remain partially true. What has changed is the documentation. The CLOUD Act clarified. Egress fees measured. The TSMC geopolitical risk identified. The absence of Capacity Planning quantified in invoices. The loss of low-level skills measurable in job postings (−34% Linux). The 2026 decision is different because the information is different — not because the decision-maker was incompetent in 2019.

Revising one's position in light of new data is not capitulation. It is the very definition of good management. That is what the Operation Dindon corpus provides: the new data. The rest is an engineering decision.

◆ WHY EGO PRECEDES EVERYTHING ELSE

The boundary between infra and dev cannot be rebuilt if the decision-maker who must authorise it cannot admit its dissolution was a mistake. All sections that follow — the risk matrix, IaC, Serverless, the sovereign IDP — are technically correct and organisationally actionable. They remain inaccessible as long as the ego lock is not lifted. That is why this section precedes all others.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2A

SECTION 2A · THE RISK AND TEMPORALITY MATRIX — A DISJUNCTION OF NATURES

NOT A SEPARATION OF POWER — A DISJUNCTION OF RISK NATURES

Before naming what was dissolved, we must name what existed — and why its dissolution is structurally dangerous.

DIMENSION	SOFTWARE DEVELOPMENT	INFRASTRUCTURE
Temporality	Short · 2-week sprint · Daily release	Long · 3-5 year contractual commitment · 10-15 year hardware horizon
Risk nature	Application bug · Functional regression	Data loss · Structural captivity · Geopolitical dependency
Reversibility	Rollback in minutes · Git revert · Feature flag	**Irreversible · Total rewrite · 6-48 months · High cost**
Primary metric	Functional velocity · Time to market · User NPS	Resilience · Uptime · TCO · Portability · Legal independence
Core competence	Application architecture · Business logic · UX · Code performance	Network/storage/compute layers · Capacity Planning · Physical security · Terms
Error consequence	Production bug · Service incident · Conversion loss	**Contractual lock-in · Loss of sovereignty · Irreversible dependency over 10 years**

◆ **WHAT THE MATRIX SAYS — THE ANSWER TO THE SILO SEMANTIC TRAP**

This table is not a defence of territory. It is a map of incompatible risk natures. A surgeon and an anaesthesiologist are not merged into the same role because they work in the same operating theatre. If the anaesthesia fails, the patient dies — regardless of the beauty of the surgical act. That is the foundation of the infrastructure engineer's independence: when they make a mistake, it is irreversible. This nature of risk demands dedicated expertise — not expertise shared with applicative velocity.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2B

THE DEVOPS MOVEMENT SOLVED A REAL PROBLEM — AND CREATED ANOTHER, DEEPER ONE

◆ WHAT DEVOPS GENUINELY ACCOMPLISHED

The DevOps movement, born in the late 2000s, responded to a documented and real problem: development and operations teams worked in watertight silos. The developer delivered an artefact "over the wall." The ops team installed it without understanding what they were installing. Time to production was measured in weeks, sometimes months. Production incidents were diplomatic crises as much as technical ones.

DevOps resolved this. Continuous integration, continuous delivery, infrastructure as code, automated testing, shared responsibility for production quality — these are real contributions that improved software quality and reduced delivery times. This study does not contest these achievements.

◆ THE UNINTENDED EFFECT — THE DISSOLUTION OF LOW-LEVEL EXPERTISE

What the DevOps movement did not anticipate: by breaking delivery silos, it also broke competence silos. The organisational merger of teams progressively led to the merger of roles, then the merger of decision-making responsibilities. The "full-stack" developer became the "full-infra" developer — not through acquired competence, but through scope drift.

Documentable result: a developer deploying on Lambda without understanding what runs underneath is not a DevOps engineer — they are someone given a remote control without being shown the television. They can change channels. They cannot repair the set. And when the set breaks down — when lock-in is installed — they do not see it coming.

The loss is double. The developer loses consciousness of low-level layers. The organisation loses the infrastructure engineer as a guard against irreversible decisions. Nobody does Capacity Planning anymore. Nobody reads Terms before choosing a managed service. Nobody asks "can this component run outside AWS without a rewrite?" — because nobody has the skills to understand the answer.

◆ THE DEVELOPER IS NOT GUILTY — THE ORGANISATION IS

The developer deploying on Lambda without understanding the infrastructure is not incompetent — they were told it was their job. The infrastructure engineer merged under an agile "tribe" is not outdated — they were told it was modernity. Both were betrayed by an organisational model that confused collaboration with the dissolution of expertise.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

YAML MADE THE HARD DRIVE FORGOTTEN — INFRASTRUCTURE-AS-CODE AS THE REAL VECTOR OF THE MERGE

The initial plan targeted the AWS console as a punitive abstraction. True — but that is the old battle. The real current vector is Infrastructure-as-Code.

◆ THE PERFECT ILLUSION — "LOOK, IT'S CODE, SO IT'S MY JOB"

Terraform, CloudFormation, Pulumi — twenty lines of HCL or YAML to instantiate a managed Kubernetes cluster (EKS), a DynamoDB database, an SQS queue, a complete VPC network. To a developer, this script looks like code. It has syntax, variables, loops, reusable modules. It lives in Git, is reviewed in Pull Requests, tested in CI/CD pipelines.

That is exactly the problem. IaC transformed the physics of infrastructure into applicative syntax. Writing `resource "aws_dynamodb_table" "users"`, the developer feels they are declaring a data structure — not signing a proprietary contractual commitment, not creating a ten-year software lock-in, not excluding any future deployment on another infrastructure without complete rewrite.

Code dissolved the consciousness of hardware. Behind twenty lines of HCL lies a datacentre in Virginia, chips fabricated by TSMC, a contract subject to California law, and egress fees patiently waiting for data to accumulate.

◆ THE DIFFERENCE THE DEVELOPER NO LONGER SEES

A BUG IN APPLICATION CODE: git revert · rollback in minutes · previous state accessible in Git history.

AN ERROR IN AN IAC SCRIPT INSTANTIATING A PROPRIETARY MANAGED SERVICE: data is inside it · application architecture depends on this service · Terraform's "git revert" does not repatriate the data · technical debt is contractually locked.

These two errors have the same syntactic form — a line of code. They do not have the same risk nature. IaC made this difference invisible. That is why the infrastructure engineer must be present in IaC script reviews — not to say no, but to name what the code represents physically and contractually.

RAITI

◆ THE ARCHITECTURAL RULE — THE MEETING POINT OF THE TWO PROFESSIONS

"Can this component run outside AWS without a rewrite?" This question, asked before each IaC resource, is the meeting point of the two professions. The developer asks it. The infra engineer answers. If the answer is no — it is an architecture decision, not a syntax choice. It must be conscious, deliberate, and documented as such.

Amine RAITI · CC BY-NC-SA 4.0

"BUILD, WE HANDLE THE REST" — THE STRATEGY TO NEUTRALISE THE INFRASTRUCTURE ENGINEER

◆ SERVERLESS AS MARKETING OF THE DISSOLUTION

The DevOps merge was not only driven by good managerial intentions. It was actively theorised, marketed and financed by hyperscalers under the banners of Serverless (AWS Lambda, 2014) and NoOps — "Build, we handle the rest." These promises have a precise target: the developer. Their declared objective: eliminate infrastructure friction. Their real effect: neutralise the infrastructure engineer as an internal opponent to cloud capture.

An infrastructure engineer who understands hardware is the primary obstacle to selling managed cloud services. They ask uncomfortable questions: What runs underneath? Where is the data? What does the contract say about egress? Can we migrate? Eliminating them organisationally — by merging their role under the developer's, by having them "evolve" toward AWS certifications, by renaming their profession "cloud engineer" — solves the commercial problem on the hyperscaler side.

◆ THE CHRONOLOGY OF NEUTRALISATION

2006-2010: the DevOps movement emerges to break delivery silos. Legitimate objective.

2012-2015: AWS, Google and Microsoft adopt and amplify the DevOps discourse in their conferences and certifications. They add the "cloud-native" layer — the idea that real DevOps happens on cloud. Hyperscaler certifications progressively replace neutral infra certifications (Linux Foundation, Red Hat).

2014-2018: Lambda (AWS), Cloud Functions (Google), Azure Functions. Serverless is presented as DevOps's next step — no more server management. The infrastructure engineer becomes progressively superfluous in this narrative.

2018-2024: organisations massively merge their infra and dev teams under DevOps "tribes." Capacity Planning disappears. Low-level skills disappear. The Linux job market collapses (−34%). Capture is complete.

The chronology is not a conspiracy theory — it is a market observation. Hyperscalers followed a rational commercial logic. Organisations followed a good-faith managerial logic. The result is documented in 64 Operation Dindon corpus studies.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHEN THE INFRASTRUCTURE ENGINEER DISAPPEARS, CAPACITY PLANNING DISAPPEARS WITH THEM — AND THE INVOICE EXPLODES

◆ THE DESTRUCTION OF CAPACITY PLANNING

The infrastructure engineer had an invisible but critical function: Capacity Planning. Before deploying a service, they estimated the load, sized the resources, anticipated peaks, calculated the monthly cost at scale. This function was not glamorous — it did not generate impressive boardroom demos. But it was the financial guardrail of the architecture.

When the infrastructure engineer was merged under the developer's umbrella, this function disappeared. The cloud engineer — often a developer retrained with an AWS Solutions Architect certification — no longer plans capacity. They consume it blindly. Lambda scales automatically — that is the promise. What the promise does not say: Lambda also scales the invoice, non-linearly, without default alerts.

◆ THE POORLY OPTIMISED LAMBDA — A QUANTIFIABLE EXAMPLE

A developer deploys a Lambda function that queries DynamoDB sequentially rather than by batch. The function is triggered 10 million times per day. Cost estimated by the developer: "Lambda is almost free at Free Tier." Real cost at scale: the combination of Lambda invocations + DynamoDB reads + data transfer can represent thousands of euros per month depending on the exact architecture.

An infrastructure engineer would have estimated this cost before deployment. Capacity planning is the translation of software architecture into real costs — before they are invoiced. Without this competence in the team, the CFO discovers costs at month-end. Not before.

◆ THE CFO AS NATURAL ALLY FOR REBUILDING THE BOUNDARY

The CFO is the only decision-maker who physically feels the consequences of the infra/dev merge — in the form of non-linearly growing cloud invoices, unforeseen exit costs, TCO never calculated. They are also the easiest to convince with data: show the cloud invoice curve over 36 months, overlay the progressive disappearance of the dedicated infra team, and let the correlation speak.

Rebuilding the infra/dev boundary is not only a technical or organisational decision — it is a financial decision. The infrastructure engineer is the CFO's guardrail. Their disappearance has a cost. This cost is measurable, documented, and growing.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3B

THE BOUNDARY IS NOT A WALL — IT IS AN ORGANISATIONAL API

◆ THE RISK OF BRUTAL RECONSTRUCTION — NOT RECREATING THE DEVS/OPS WAR

Recreating a watertight boundary between infra and dev without a protocol immediately reactivates the legendary "Devs vs Ops" war: the JIRA ticket sleeping for three weeks, the infra team reflexively saying no, the dev team circumventing out of frustration. This war existed. It produced the DevOps movement. The goal is not to return to 2005.

The boundary to rebuild is not a wall. It is an interface — an organisational API that allows the two professions to collaborate without dissolving into each other.

◆ THE SOVEREIGN IDP — THE INFRASTRUCTURE ENGINEER AS PLATFORM PROVIDER

The sovereign Internal Developer Platform (IDP) is the operational answer. The infrastructure engineer no longer says "no" — they provide. They design and make available to developers abstract infrastructure bricks that are agnostic and reversible: a database brick that can run on PostgreSQL bare-metal or on cloud, a message queue brick that can be RabbitMQ or SQS depending on context, an authentication brick not locked to Cognito.

The sovereign brick rule: every component provided by the internal platform must be able to answer "yes" to the question "can it run outside AWS without a rewrite?" If the answer is no, the choice is deliberate, documented, and assumed by both parties — not suffered by default because "it was easy to deploy."

◆ EPISTEMOLOGICAL TRAINING — THE MATHS TEACHER'S BAG FOR INFRA

The IDP solves the organisational problem. It does not solve the epistemological problem: a developer using an IDP brick without understanding what is underneath remains unable to recognise an architectural lock-in when they create one.

The solution is not to train every developer to manage a Kubernetes cluster. It is to show them once, physically, what `resource "aws_dynamodb_table"` triggers underneath — like the teacher who takes out a satellite dish so the student understands sine and cosine. Not to make them an infrastructure engineer. So they can read a cloud invoice, recognise an inode, understand why an IaC script is not a git revert.

HUMAN

◆◆◆

The Iron does not lie.

It has no Free Tier. It has no YAML console.

It costs from day one — and that is exactly why it is free.

Amine RAITI · Operation Dindon · 2026

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026

THE CLOUD-WASHING

Anatomy of an Imposture in Three Layers

Sovereignty Washing · Green Washing · Finance Washing

HUMAN

◆ THE THESIS

Cloud-Washing is the meta-mechanism that overarches all others: it paints dependency as freedom, domination as service, extraterritoriality as sovereignty. It operates in three distinct layers — sovereignty, ecology, and public finance — each with its own rhetorical tools, its own labels, and its own figures. This study maps them. The poem that opens the document named them before the analysis documented them.

LAYERS	FORMAT	WATERMARK
3	**POEM + RATIO**	**HUMAN + RATIO**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

MU'ALLAQA IV

THE CLOUD-WASHING | STYLE: AL-MUTANABBI

This poem was written by Amine RAITI before this study. The study documents what the poem named.

Cursed be the ruse and the paint of the great, · Who hide their empire in words of high state.
Azure and AWS, these masters of guile, · Paint Sovereign on every vile, tired mile.
They speak of the green and the sky's open grace, · While Microsoft remains as the mask on the face.
'Tis a visor of iron on a face made of clay, · A marketing court for the weak of the day.
Woe to the falsehoods in the colors of old, · Where Google puts us to sleep with the lies they have told.
Cloud-Washing is there, like a veil made of silk, · To hide that Amazon owns the cream and the milk.
The engineer blinded by labels of straw, · Gives up his keys to a foreigner's law.
They sell us the homeland with an alien seal, · To ensure that no one can break through the steel.
Woe to the finance that pays for the show, · In the Cloud of the Empire, where the vanities grow.
We pay for the luxury, for a local veneer, · While Microsoft reaps a gain that is clear.
The gold of our taxes feeds these giants of old, · Who laugh at our laws, so tiny and cold.
'Tis a tax that is hidden on the mind's own light, · A levy on negligence in the dead of the night.
The poet arises and denounces the king, · Who knows only the law that his own masters bring.
I reject these words and these predators' breed, · These merchants of Google, these sowers of greed.
True power belongs to the one who possesses, · The physical machine that no stranger oppresses.
The Bare-Metal laughs at your labels of card, · For the Iron lies not, 'tis our staff and our guard.
Let the Expert beware of the velvet discourse: · The Cloud is no roof, but a dark, hollow source.
Mistake not the label for the truth of the tree, · Or lose very soon all your right to be free.

HUMAN

AMINE RAITI · OPERATION DINDON · JUNE 2026 · CC BY-NC-SA 4.0

I

SECTION I · SOVEREIGNTY WASHING

"SOVEREIGN" ON INFRASTRUCTURE SUBJECT TO CALIFORNIA LAW

Azure and AWS, these masters of guile, / Paint Sovereign on every vile, tired mile.

◆ THE HYPERSCALERS' SOVEREIGN OFFERINGS

Since 2022-2023, all three hyperscalers have launched sovereign offerings targeting the European market:

AWS European Sovereign Cloud: physically in Europe, operated by European employees, with enhanced access controls. What the offering does not say: Amazon.com Inc. remains subject to the US CLOUD Act. A National Security Letter can compel Amazon to provide data without informing the client, regardless of server location.

Microsoft Cloud de Confiance with Bleu: a French subsidiary (Bleu, Capgemini + Orange) operates Microsoft infrastructure in France. Technology remains Microsoft's. Contract remains subject to Irish and US law for the underlying technology.

Google Sovereign Controls: client-managed encryption reducing Google's data access. Partial reduction of exposure — not complete sovereignty.

◆ THE DIFFERENCE THE LABEL DOES NOT SAY

Physical sovereignty: where are the servers? Resolved. Servers are in Europe.

Legal sovereignty: which law applies? Not resolved. The parent company remains American, subject to the CLOUD Act and California law (Terms §14.12 documented in Terms Under the Microscope).

Sovereignty Washing sells the physical resolution as if it also resolved the legal dimension. The decision-maker who reads Sovereign Cloud believes both are resolved. Only one is.

RATIO

◆ THE VERSE

They sell us the homeland with an alien seal, / To ensure that no one can break through the steel.

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · GREEN WASHING

THE VIRGINIA DATACENTRE AND ITS REAL ENERGY MIX

They speak of the green and the sky's open grace, / While Microsoft remains as the mask on the face.

◆ RECS · PPAS · AND THE REALITY OF THE ELECTRICAL GRID

RECs (Renewable Energy Certificates): a certificate attesting that an equivalent quantity of energy was injected into the grid by a renewable source. The server hosting your data continues consuming whatever energy is available locally — which may be predominantly fossil. RECs balance the accounting, not the grid.

PPAs (Power Purchase Agreements): long-term renewable energy purchase contracts, more substantive than RECs. But a PPA signed in Spain does not mean the Virginia datacentre runs on Spanish wind. Electrons do not travel that way.

The Northern Virginia electrical grid — where the largest AWS and Microsoft datacentres concentrate — is powered approximately 30-35% by fossil sources according to 2023 EIA (US Energy Information Administration) data.

◆ THE HIDDEN INEFFICIENCY OF CLOUD ARCHITECTURES

Pooling genuinely reduces consumption versus individual under-utilised datacentres. But poorly designed serverless architectures — Lambda cold-starting, idle containers, monitoring querying APIs every 30 seconds — consume compute continuously. Nobody measures kWh per request. The green pooling is real at datacentre scale — cancelled by the application inefficiency cloud encourages through its billing model.

◆ THE VERSE

'Tis a visor of iron on a face made of clay, / A marketing court for the weak of the day.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · FINANCE WASHING

PUBLIC MONEY FINANCING THE DEPENDENCY IT CLAIMS TO COMBAT

The gold of our taxes feeds these giants of old, / Who laugh at our laws, so tiny and cold.

◆ PUBLIC CLOUD PROCUREMENT IN FRANCE AND EUROPE

While European governments produce reports on digital sovereignty, public administrations sign cloud contracts with the same hyperscalers they claim to want to reduce.

In France: DINUM selected AWS, Azure and Google in its Cloud au Centre programme for public administrations. Public hospitals migrate to Microsoft Azure. Universities use Google Workspace for Education. Ministries deploy on AWS.

The European Union itself massively uses US hyperscaler services for its own institutions — while funding the GAIA-X programme intended to create a European alternative.

The Finance Washing paradox: public money funds hyperscalers through procurement. These same hyperscalers use part of this money for lobbying on European regulations (GDPR, Data Act, AI Act). European taxpayer money indirectly funds resistance to the regulations that would protect them.

◆ PUBLIC TRAINING FUNDS FINANCING PROPRIETARY CERTIFICATIONS

Documented in The CPF Misused: cloud certification spending tripled between 2020 and 2023. Public professional training money massively funds AWS, GCP and Azure certifications. These certifications reinforce dependency by creating a labour market that values proprietary skills above neutral skills. The public training system funds the cognitive capture documented in The Newspeak That Costs Dear.

◆ THE VERSE

'Tis a tax that is hidden on the mind's own light, / A levy on negligence in the dead of the night.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · TABLE OF IMPOSTURES — THREE LAYERS MAPPED

WHAT THEY ANNOUNCE · REALITY · THE COST IN SOVEREIGNTY

LAYER	THEY ANNOUNCE	REALITY	COST IN SOVEREIGNTY
SOVEREIGNTY WASHING	*Sovereign Cloud · Trusted Cloud · Sovereign Controls*	Servers Europe (physical OK) · California jurisdiction (legal NO) · CLOUD Act applies	Decision on false premise · Legal dependency maintained · Exit made harder
GREEN WASHING	*100% renewable · Net zero carbon · Green cloud*	RECs not equal physical green energy · Virginia 30-35% fossil · kWh/request unmeasured	Sustainability team validates migration on ecological argument · Incomplete ESG
FINANCE WASHING	*Public service modernisation · Digital transformation · Cloud au Centre*	Public money hyperscalers lobbying anti-regulation · Training funds proprietary certifications	Taxpayer funds their own capture · Self-sustaining loop

◆ CLOUD-WASHING AS META-MECHANISM

The three layers neutralise the three internal resistances: Sovereignty Washing reassures legal · Green Washing reassures sustainability · Finance Washing reassures CFO (OpEx vs CapEx). The engineer who says let us calculate the real TCO stands alone against three arguments and three directorates. Cloud-Washing is not an isolated lie — it is an architecture of conviction.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · LINK WITH THE CORPUS

THE POEM AS STARTING POINT · THE CORPUS AS DOCUMENTATION

The poem MU'ALLAQA IV was written before this study. It named the three layers with the precision of direct observation. Al-Mutanabbi denounced the powerful at their own court. The Operation Dindon corpus does the same — in verse and comparative tables.

◆ WHAT THE CORPUS HAD ALREADY DOCUMENTED

Sovereignty Washing: Data residency as false friend · Compliance shifting responsibility · Global infrastructure masking extraterritoriality. Documented in The Newspeak That Costs Dear.

Green Washing: RECs not equal physical green energy · kWh/request unmeasured · Sustainability team ticking the box. Documented in The Newspeak blind spots (Section 7 GreenOps).

Finance Washing: training funds proprietary certifications · public money cognitive capture. Documented in The CPF Misused.

◆ WHAT THIS STUDY ADDS

The term Cloud-Washing unifies these three mechanisms under a common label and says they are not three independent problems — they are three layers of the same conviction system designed to validate migration through three different directorates with three different arguments. The engineer who would say no must now convince all three simultaneously.

◆◆◆

The Bare-Metal laughs at your labels of card,

For the Iron lies not, 'tis our staff and our guard.

Amine RAITI · MU'ALLAQA IV · 2026

◆◆◆

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026

THE FREE SAMPLE

Mechanics of Dependency Applied to Cloud

Why Free Tier follows the logic of every industry that discovered how to create a need before monetising it

HUMAN

◆ THE THESIS

Free Tier is an entry sample. Not a commercial offer — a sample. The food retail, video game, and pharmaceutical industries all discovered the same principle: offer the first experience free to create the need, then monetise the impossibility of going without. AWS applied this mechanic to digital infrastructure with surgical precision. This study draws the structural parallel — without a single hard word — and leaves the reader to draw the conclusion.

NEUTRAL TERMS	REGISTER	WATERMARK
10	**SUBTLE NO HARD WORDS**	**HUMAN**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

I

SECTION I · THE QUESTION THAT OPENS EVERYTHING

WHICH INDUSTRY INVENTED THE PRINCIPLE OF GIVING FREE TO CREATE A NEED THAT WOULD THEN BE MONETISED?

◆ THE SAMPLE IN INDUSTRIES THAT DISCOVERED THE MECHANIC

FOOD RETAIL: supermarket samples. Cheese cut in small cubes. Sausage on a cocktail stick. Free. No commitment. Just to taste. The food industry discovered in the 20th century that physical product experience converts better than any advertising. Once you have tasted, the need exists. It did not exist before.

VIDEO GAME INDUSTRY: the free-to-play model. The game is free. Progression is free. Until the game is designed to create a frustration only a purchase resolves. Loot boxes, slowly recharging "energy," levels blocked before a paywall — same principle: free sample, paid friction. The European Union began regulating these practices in 2022.

PHARMACEUTICAL INDUSTRY: sample medicines distributed to doctors. Free for the patient in the first months. When the prescription follows, the molecule is patented and the generic does not yet exist. The medical need is real — but it was created by free initial access.

AWS · GCP · AZURE: Free Tier. Startup Credit. AWS Educate. Free Ingress. Same principle. Same mechanic. Same economic structure. The only difference: the product is not food, entertainment, or medical. It is digital infrastructure. But the sequence is identical: free sample → creation of need → functional dependency → monetisation of the impossibility of exit.

◆ WHAT THE PARALLEL REVEALS

These industries were all regulated when the public understood the mechanic. Food retail has rules on misleading promotions. Video gaming faces regulatory pressure on loot boxes. The pharmaceutical industry is under surveillance on sample conflicts of interest. The cloud industry has not yet been regulated on this specific dimension. The ultimatum of 8 September 2026 is a first step in that direction.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE CORRESPONDENCE TABLE — NEUTRAL TERMINOLOGY AND CLOUD EQUIVALENTS

TEN TERMS · TEN EQUIVALENTS · ONE MECHANIC

The table below uses strictly neutral terminology — drawn from economics, marketing, and behavioural psychology. Not a single term is borrowed from a pejorative register. The conclusion belongs to the reader.

NEUTRAL TERM	DEFINITION	CLOUD EQUIVALENT
The entry sample	Product offered free to create the first contact	*Free Tier · AWS Educate · Startup Credits*
The integration substance	What creates functional dependency once adopted	*Proprietary APIs · DynamoDB · Lambda · S3*
Functional dependency	State where stopping causes unbearable friction	*Software lock-in · Code written for a single vendor*
The tipping threshold	Moment where the free stops and billing begins	*End of 12 months · Startup credits exhausted*
Exit friction	High cost and pain associated with stopping consumption	*Egress fees · TCO crisis · Forced migration*
The captive ecosystem	Network that makes exit socially and technically difficult	*Proprietary certifications · AWS labour market*
The proximity distributor	First-contact agent who distributes samples	*AWS commercial on campus · Accelerator programmes*
The ecosystem architect	The one who designs and controls the entire chain	*AWS · GCP · Azure — the hyperscalers*
Chosen withdrawal	Return to independence — difficult but possible	*Bare-metal migration · Accompanied Exodus*
The product without sample	What costs from day one and needs no dependency creation	*Bare-Metal · Owned physical server*

◆ THE FOUR-STEP SEQUENCE — IDENTICAL IN ALL THESE INDUSTRIES

1. Free entry sample — first contact with no financial friction. 2. Functional dependency — the product integrates into habits, workflows, architectures. 3. Tipping threshold — the free stops, billing begins, exit has become costly. 4. Exit friction — the captive ecosystem makes return to independence long, costly, and risky.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE PRODUCT WITHOUT A SAMPLE — THE COUNTER-PROPOSAL

BARE-METAL COSTS FROM DAY ONE · THAT IS ITS MAIN ADVANTAGE

◆ WHY THE ABSENCE OF A SAMPLE IS A PROPERTY, NOT A DEFECT

Bare-Metal has no Free Tier. No Startup Credit. No free Ingress. It costs from day one — electricity, rack, maintenance, competent staff. This absence of a sample is often cited as a disadvantage: "cloud is easier to start." Correct. And that is precisely why Bare-Metal creates no functional dependency. A product that costs from day one has no need to make you dependent — you have already made your decision in full knowledge of cost. It has nothing to offer you for free. It does not need to push you through a tipping threshold. It has no exit friction because there is no exit — you own the hardware.

◆ THE PARALLEL WITH INDUSTRIES THAT RESISTED THE SAMPLE MECHANIC

Not all industries adopted the sample mechanic. The craftsman invoices from the first cut of the chisel. The architect charges fees from the first sketch. The independent consultant sends an invoice from the first call. These professionals have no Free Tier because they do not need one — their value is visible without a free sample.

The sovereign managed service provider who offers services to a client is in the same position: it invoices from day one. It cannot offer 12 free months because it does not have Amazon's financial reserves. But this constraint is also its freedom: it does not need to create dependency to retain clients. It retains them because it serves them well.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · BLIND SPOT · THE PRIDE LOCK

THE SUNK COST FALLACY APPLIED TO CLOUD — LOCK-IN IS ALSO AN ARCHITECTURE OF EGO

The corpus documented all the technical and financial barriers to exit. It had not yet documented the most powerful lock — the one that operates at the exact level where decisions are made.

◆ THE SUNK COST FALLACY — DEFINITION AND MECHANIC

The Sunk Cost Fallacy (Thaler 1980, Kahneman 2011) is one of the most robust and universal cognitive biases: the tendency to continue a commitment not because it is rational to continue, but because one has already invested non-recoverable resources. Spectators watch a bad film to the end because they paid for the ticket. States continue lost wars because they have already invested soldiers. CIOs continue paying egress fees because they have already invested five years of their career in the cloud migration.

This bias is particularly powerful when past investment is public and documented. And cloud migrations are public — they are the subject of LinkedIn posts, conferences, use cases published on the vendor's website, internal testimonials. The CIO who evangelised cloud for five years has not only invested time and money. They have invested their professional reputation.

◆ THE PROFILE OF THE TRAPPED DECISION-MAKER

The CIO concerned has typically: published LinkedIn posts on the successful migration · presented the use case at conference · received a promotion or internal recognition linked to the digital transformation · convinced their CEO and board that cloud was the right decision · trained their teams on AWS/GCP certifications.

Five years later, they read the Dindon corpus. The figures are correct. The 7.5x bare-metal vs cloud ratio is verifiable. Terms §14.12 are in their contracts. They understand. And they cannot retreat — not because they cannot calculate the TCO, but because retreating means publicly admitting that five years of LinkedIn posts were a mistake. The fear of losing face before their CEO and peers is stronger than economic rationality. This is sunk cost operating at its maximum.

◆ THE PROPOSAL — HOW TO LET A CIO RETREAT WITHOUT LOSING FACE

Migration toward bare-metal is not an admission of error — it is a strategy evolution. The context has changed (CLOUD Act clarified, egress fees documented, TSMC geopolitical risk identified). The 2019 decision was correct with 2019 information. The 2026 decision is different because the information is different. This is good management — not capitulation.

This narrative is not communication — it is the truth. The Dindon corpus documents mechanisms that were not as clearly articulated in 2019. The CIO who revises their position in 2026 in light of the corpus is no more wrong than the engineer who revises their calculations when better data is available.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · BLIND SPOT · ARTIFICIAL COMPLEXITY — THE CHEMICAL AGENT OF CAPTIVITY

A PERFUME SAMPLE DOES NOT MODIFY YOUR BIOLOGY · A CLOUD SAMPLE MODIFIES THE ARCHITECTURE OF WHAT YOU BUILD

This is the fundamental difference between cloud and all other sample industries — and what makes cloud captivity structurally deeper.

◆ THE LIMIT OF THE SAMPLE METAPHOR

In all industries documented in Section 1, the sample creates functional dependency — but does not modify the consumer's structure. A cheese sample creates a desire for cheese. It does not transform your digestive system to be able only to ingest that precise cheese. A sample medicine creates a treatment need. It does not reconfigure your biology to make your body incompatible with other molecules.

The cloud sample does something fundamentally different: it modifies the very structure of the product you build. When a developer builds on AWS Free Tier with Lambda + DynamoDB + API Gateway + SQS + Cognito + CloudFront, they are not building an application that runs on AWS. They are building an application that **CAN ONLY RUN ON AWS**. The architecture becomes biologically incompatible with any other environment.

◆ ARTIFICIAL COMPLEXITY AS A CAPTURE MECHANISM

Each nested managed service is an additional atom of captivity. Lambda triggering SQS triggering another Lambda writing to DynamoDB with a Global Secondary Index read by an API Gateway with a Cognito authorizer — each link in this chain is a service that does not exist outside AWS in this exact form. The entire chain is proprietary. Migrating this system to bare-metal or another vendor is not a migration — it is a total rewrite.

Cloud-Washing (MU'ALLAQA IV) presents this complexity as modernity. The terms "serverless," "event-driven," "fully managed," "cloud-native" are modernity labels applied to what is in reality artificial complexity — an architecture designed not to solve a technical problem but to make exit impossible.

◆ THE QUESTION THE ARCHITECT MUST ASK BEFORE EACH MANAGED SERVICE

"CAN THIS COMPONENT RUN OUTSIDE AWS WITHOUT A REWRITE?"

If the answer is no — it is artificial complexity, not modernity. The question does not condemn using the service. It requires a conscious decision: I choose this service knowing it ties me to AWS. That is a sovereign decision. What Free Tier prevents is this decision being made consciously — because free-ness anaesthetises judgment on real costs, including the cost of architectural captivity.

The developer who asks this question before each managed service does what The Maths Teacher's Bag teaches the student: see the physical reality behind the abstract symbol. The managed service is the symbol. Architectural captivity is the physical reality.

◆◆◆

Bare-Metal is expensive from day one.

That is its main advantage.

It has nothing to offer you for free.

It does not need to make you dependent.

It belongs to you.

Amine RAITI · Operation Dindon · 2026

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026

THE FREE TIER BAIT

Capturing the Developer Before the Migration

7 Mechanisms · From Free to Ransom

◆ THE THESIS

Before the Terms. Before the newspeak. Before the boardroom. There is the Free Tier. This study documents seven mechanisms of the Free Tier Bait: student capture, the invisible threshold and brutal invoice, Free Tier in public procurement, the startup credit as financial anaesthesia, the data gravity trap, the State's budgetary abandonment, and the complete causal chain. Nothing is free. Entry is free because it creates dependency. Exit is charged because it reveals its cost.

MECHANISMS	PAGES	FORMAT	WATERMARK
7	10	POEM + RATIO	HUMAN + RATIO

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

HUMAN

MU'ALLAQA IX

The Free Tier Bait | Style: Zohayr

This poem was written by Amine RAITI before this study. The study documents what the poem named.

Cursed be the offer where the gift is a snare, . To capture the mind in the first morning's glare.
AWS lures us with the gifts that they bring, . Hiding the burden 'neath a free-flowing thing.
Google offers access, Microsoft invites, . To taste of the honey that poisons our nights.
'Tis a war that is slow with the face of a friend, . Where we give up our keys for a bitter-sweet end.
Woe to the limit that we never see near, . Which makes of the free but a memory of fear.
The student falls asleep on his code, light and thin, . But Azure in secret is counting the sin.
One day the bill falls, brutal and bare, . For the threshold is crossed and the trap's in the snare.
We master no more, we but suffer the rate, . Of a cruel, dark system that seals our own fate.
Woe to the finance that sows just to reap, . The blood of the innocent while they're asleep.
Amazon watches and calculates gain, . While we build our world on their own cold terrain.
The free-tier shines bright to make us all slaves, . Of these Hyperscalers and their digital graves.
'Tis a usury vile on the knowledge of old, . A tax on the dream, so icy and cold.
Zohayr turns away and refuses the gift, . Which is the only true cause of our spiritual rift.
I reject these offers and these benefactors' pride, . These merchants of Free Tier with nowhere to hide.
The true master knows that nothing is free, . In this foreign Cloud where we bend on our knee.
The Bare-Metal costs, but it never betrays, . For it waits not the hour of your wandering ways.
Let the Expert remember the hook in the gold: . The gift of the giants is a story of old.
Mistake not the trial for a freedom so vast, . Or pay for the price while the shadows are cast.

HUMAN

Amine RAITI · Operation Dindon · Juin 2026 · CC BY-NC-SA 4.0

CC BY-NC-SA 4.0

I

SECTION I · THE STUDENT CAPTURE

AWS EDUCATE · GOOGLE FOR EDUCATION · AZURE FOR STUDENTS — ACQUISITION AT ZERO MARGINAL COST

The student falls asleep on his code, light and thin, / But Azure in secret is counting the sin.

◆ THE HYPERSCALERS' EDUCATIONAL PROGRAMMES

AWS Educate gives students and teachers access to cloud credits and training. Google for Education provides Google Workspace free to millions of institutions. Microsoft Azure for Students offers $100 in credits with no credit card required. These programmes reach tens of millions of students worldwide — including future infrastructure architects, future CIOs, future technical decision-makers.

This is not philanthropy. It is customer acquisition at the lowest possible marginal cost. A student who learns to code on AWS learns AWS APIs, AWS conventions, AWS services. Their reflexes, code snippets, portfolio projects, certifications — all AWS. When they join a company five years later and recommend a tech stack, they recommend what they know. What they know is what was given to them for free.

◆ THE SCHOLASTIC AMPUTATION REINFORCED BY FREE TIER

The study Sir, What Is This Good For documented that school creates a rupture between symbol and matter. Free Tier adds another layer: it gives access to cloud abstractions without the student ever having configured a physical server, opened a network cabinet, or manipulated a real filesystem. The developer trained on Free Tier knows how to deploy a Lambda function. They do not know what runs underneath. They have never needed to — Free Tier removed that necessity.

◆ THE VERSE

'Tis a usury vile on the knowledge of old, / A tax on the dream, so icy and cold.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE INVISIBLE THRESHOLD AND THE BRUTAL INVOICE

DESIGNED TO BE EXCEEDED WHEN ATTACHMENT IS MAXIMAL AND EXIT MOST COSTLY

One day the bill falls, brutal and bare, / For the threshold is crossed and the trap's in the snare.

◆ THE EXACT FREE TIER LIMITS — AWS · GCP · AZURE

AWS Free Tier: 750h EC2 t2.micro/month for 12 months · 5 GB S3 · 1 million Lambda requests/month · 25 GB DynamoDB. Documented — but no real-time alert in the developer's workflow. No red banner in the AWS console when approaching the threshold.

GCP Free Tier: 1 e2-micro instance/month · 5 GB Cloud Storage · 2 million Cloud Functions requests. Same principle: limits invisible in the workflow.

Azure Free Tier: 750h B1S virtual machines for 12 months · 5 GB Blob Storage · 1 million Functions requests.

◆ THE BILL SHOCK MECHANISM — NOT AN ACCIDENT

The project grows. Users arrive. Traffic rises. The developer codes — not monitoring their invoice. Then: first invoice. Brutal. At that point, migrating costs more than paying. Code is in DynamoDB. Infrastructure is in CloudFormation. Developers know AWS — nothing else.

Free Tier is designed so the threshold is crossed exactly when attachment is maximal and exit is hardest. Not a design accident. The design.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · FREE TIER IN PUBLIC PROCUREMENT

THE PUBLIC BODY THAT TESTS FOR FREE AND PAYS INDEFINITELY WITH TAXPAYER MONEY

Amazon watches and calculates gain, / While we build our world on their own cold terrain.

◆ THE TYPICAL SCENARIO IN PUBLIC ADMINISTRATION

A local authority IT department wants to modernise a public service. Budget constrained. A contractor proposes: "let's test on cloud — it's free to start." Project begins on AWS or Google Free Tier. Six months later: threshold exceeded. The authority is a paying customer of a US hyperscaler for a public service funded by local tax. No procurement process — the project started free and did not reach public contract thresholds. The contractor is gone. Internal migration skills do not exist.

◆ THE FINANCE WASHING x FREE TIER PARADOX

Public Free Tier is the junction of two mechanisms documented separately. On one side, Finance Washing (MU'ALLAQA IV): public money funds hyperscalers through procurement. On the other, the Free Tier Bait: project starts free, threshold crossed when attachment is maximal. In the private sector, the decision-maker commits their company. In public, they commit taxpayer money — often without formal procedures. The taxpayer pays twice: once for the public service, once for an exit they cannot fund.

◆ THE VERSE

Woe to the finance that sows just to reap, / The blood of the innocent while they're asleep.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · BLIND SPOT · THE STARTUP CREDIT — THE GRANT OF FINANCIAL ANAESTHESIA

AWS ACTIVATE · GOOGLE FOR STARTUPS · MICROSOFT FOR STARTUPS — DISGUISED VENTURE CAPITAL

The free-tier shines bright to make us all slaves, / Of these Hyperscalers and their digital graves.

◆ THE STARTUP CREDIT MECHANISM

Beyond the consumer Free Tier, hyperscalers distribute massive cloud credit packages to early-stage startups: AWS Activate up to $100,000 · Google for Startups up to $200,000 · Microsoft for Startups up to $150,000. These credits are distributed through accelerators, incubators, and partner venture capital funds.

This is not Free Tier. It is disguised venture capital. The effect is radically different: a startup with $100,000 in cloud credits does not build frugally — it builds like a CAC40 firm. Microservices, proprietary PaaS (managed Kubernetes, database-as-a-service, managed AI), complex distributed architectures. Since it pays nothing, it does not optimise. It accumulates cloud technical debt without knowing it.

◆ THE TCO CRISIS — A NEW TERM FOR THE CORPUS

When credits run out — generally between 12 and 18 months — the startup discovers its true infrastructure cost. The "TCO crisis" is the moment when the real cost structure imposes itself brutally: an architecture built for free capital cannot be paid for with the real revenues of a growth-stage startup. Two outcomes: raise additional funds to pay infrastructure (entering additional financial dependency on VCs) or migrate under emergency conditions — taking months and risking breaking everything.

VC funds push to consume credits rapidly to accelerate "Time to Market." They have signed partnerships with hyperscalers — each startup in their portfolio consuming credits is a new acquisition for AWS. The VC, the startup, and the hyperscaler form a triangle whose first two fund the third without knowing it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · BLIND SPOT · THE DATA GRAVITY TRAP — INGRESS FREE · EGRESS CHARGED

ENTRY IS FREE BECAUSE IT CREATES DEPENDENCY · EXIT IS CHARGED BECAUSE IT REVEALS ITS COST

We master no more, we but suffer the rate, / Of a cruel, dark system that seals our own fate.

◆ THE PERMANENT FREE TIER ON DATA INGRESS

Hyperscalers offer data transfer into their cloud (Ingress) completely free — permanently, not limited to 12 months. Sending 100 TB into AWS S3: free. Sending 1 PB into GCP Cloud Storage: free. This is an endless Free Tier on data ingress.

The illusion of free works. The organisation migrates its data without financial friction. The decision seems rational — no immediate cost, no visible risk. The data enters. The data accumulates. Metadata is created. Indexes are built. Applications connect. Dependency installs itself silently.

◆ EGRESS FEES — THE EXIT RANSOM

When the organisation wants to exit or repatriate data to its physical hardware, pricing reverses brutally:

AWS: $0.09/GB for the first 10 TB outbound per month (to internet) · regional rates vary · free to other AWS services (creating additional inter-service dependency).

GCP: $0.08/GB for internet egress · free internally.

AZURE: $0.087/GB for first 10 TB outbound.

On 100 TB of data: $8,000 to $9,000 in egress fees. On 1 PB: $80,000 to $90,000. This is not a service fee — it is an exit ransom. It was documented in Terms Under the Microscope as a contractual retention mechanism. It is also a reverse Free Tier: free entry funds the paid exit.

◆ THE FORMULA

Entry is free because it creates dependency. Exit is charged because it reveals its cost. The data gravity trap is the most durable Free Tier — it does not last 12 months, it lasts as long as the data remains in the cloud.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · BLIND SPOT · THE STATE'S BUDGETARY ABANDONMENT — WHY THE TEACHER HAS NO CHOICE

AWS EDUCATE EXISTS NOT BECAUSE AMAZON IS PHILANTHROPIC · BUT BECAUSE THE STATE ABANDONED PEDAGOGICAL INFRASTRUCTURE

Woe to the limit that we never see near, / Which makes of the free but a memory of fear.

◆ THE COMPUTING TEACHER'S BINARY CHOICE

A computing teacher in a public university wants students to handle real workloads — servers, networks, databases, deployments. Two options.

Option A: use free consoles provided by AWS, Google or Microsoft. Students handle modern, real tools in functional environments. The course is credible. Students learn — vendor APIs, but they learn something.

Option B: not use proprietary consoles. In this case, students handle no modern tools. The university has no budget for lab servers, no room dedicated to physical experimentation, no cluster available for student projects.

The teacher does not choose between AWS and a sovereign alternative. They choose between AWS and nothing. Not negligence — institutional budget constraint.

◆ THE VOID HYPERSCALERS IDENTIFIED AND FILLED

AWS Educate does not exist out of philanthropy. It exists because Amazon identified that universities had no server budget — and that this void was an opportunity to acquire future captive clients. The State progressively eliminated physical pedagogical infrastructure budgets. Computer lab rooms with real servers disappeared. Computing clusters for student projects are rare. The void exists. Hyperscalers filled it with Free Tier.

Not a coincidence. A strategy. The solution documented in The Maths Teacher's Bag (satellite dish, screws, marble in the parabolic dish) applies here too: not a programme reform, but a physical server in every university. Not an AWS console — a server the student can touch, configure, and break without additional billing.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE COMPLETE CAUSAL CHAIN — THREE ENTRY DOORS · ONE EXIT

FROM THE FIRST AWS CONFIGURE TO THE ULTIMATUM OF 8 SEPTEMBER 2026

THREE ENTRY DOORS

1. Student Free Tier
Student → proprietary APIs → no fundamentals → AWS reflexes → recommends AWS at work

2. Startup Credit
Startup → $100k credits → complex architecture → credits gone → TCO crisis → captive or dead

3. Free Ingress
Data enters free → accumulation → application dependency → egress charged → exit ransom

THE COUNTER-CHAIN — BARE-METAL FROM DAY ONE

One physical server in every university
↓ Student configures, breaks, starts again
↓ Fundamentals acquired (Linux, network, storage)
↓ No Free Tier — no threshold — no bill shock
↓ Neutral exportable skills
↓ Startup sized to real revenues
↓ Data on own hardware — egress = zero
↓ Local authority masters its infrastructure
↓ **Ultimatum possible · Dindon corpus as documentation**

◆ WHAT THE SEVEN MECHANISMS HAVE IN COMMON

Student Free Tier · Invisible threshold · Public procurement · Startup credit · Data gravity trap · Budgetary abandonment · Causal chain — seven mechanisms, one logic: make cloud free or nearly free when dependency installs, and expensive or very expensive when exit is considered. Free Tier is not a commercial offer. It is an architecture of capture.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

CLOSING · FREE TIER IS AN ARCHITECTURE OF CAPTURE

Free Tier never truly ends. It evolves. It grows. It changes form. The student Free Tier becomes the Startup Credit. The Startup Credit becomes free Ingress. Free Ingress becomes paid Egress. Initial free becomes permanent dependency.

Meanwhile, the State has no budget for a server in every university. The teacher opens the AWS console for lack of anything better. The student learns vendor APIs. The startup builds on free capital. The local authority signs without procurement. And five years later, the CIO signs Terms §14.12.

The chain is complete. It began with a click on "Start your free trial."

The Bare-Metal costs, but it never betrays,
For it waits not the hour of your wandering ways.

Amine RAITI · MU'ALLAQA IX · 2026

NEMO SUPRA LEGEM EST

HUMAN

CC BY-NC-SA 4.0

SARCASTIC STUDY · OPÉRATION DINDON · JUNE 2026

THE €3,000 VM AND THE €2,800 SRE

When the Machine Earns More Than the Engineer Who Keeps It Running

A Study in Creative Accounting

◆ THE TWO NUMBERS THAT SAY EVERYTHING

A GCP n2-highmem-32 instance (32 vCPU, 256 GB RAM) costs approximately €3,000 per month on-demand. The SRE who manages it earns approximately €2,800 per month net. The machine earns more than the engineer. The organisation spends more renting a machine it will never own than paying the human who keeps it running 24/7. And when budgets are tight, guess which one goes first.

VM / MONTH	SRE / MONTH NET	WATERMARK
€3,000	**€2,800**	**HUMAN**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

HUMAN

I

SECTION I · THE CALCULATION NOBODY DOES IN THE BOARDROOM

5 YEARS OF VM = €180,000 · 5 YEARS OF BARE-METAL = €24,000 · RATIO: 7.5×

The board looks at the "cloud infrastructure" line and sees modernity, flexibility, clean OpEx. It does not see €180,000 over 5 years for a machine it will never own. If the CFO saw the same figure next to "apartment rent", they would immediately ask whether it is possible to buy.

ITEM	GCP VM ON-DEMAND	NEW BARE-METAL	REFURBISHED GRADE A
Acquisition cost	€0 (but you own nothing)	€15,000 (you are the owner)	€5,000 (you are the owner)
Monthly cost	€3,000/month · forever	€250/month (depreciation) + €150 opex = €400	€140/month (depreciation) + €130 opex = €270
5-year total cost	**€180,000**	**€24,000**	**€16,200**
Residual value at 5 years	€0 — you own nothing	€2,000-5,000 (resale or spare)	€500-1,500 (resale)
Ratio vs VM	**1× (reference)**	**7.5× cheaper**	**11× cheaper**

◆ WHAT THE TABLE DOES NOT YET SHOW

The €3,000/month GCP VM does not include egress fees when you try to retrieve your data, inter-service network fees, snapshot fees, load balancer fees, monitoring fees, Premium support fees if you want to talk to a human. The real ratio between the all-in GCP VM and bare-metal is often 10× to 15×. And in 5 years, if GCP decides to raise prices "at any time" (Terms of Service Section 2.6), the meter goes up again. You have nothing to renegotiate — because you own nothing.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE MACHINE THAT EARNS MORE THAN THE ENGINEER

€3,000/MONTH THE VM · €2,800/MONTH THE SRE · AND THE SRE IS "A COST"

In the organisation's accounting, the €3,000/month VM appears as an "infrastructure cost" — modern, justifiable, in the right Excel spreadsheet under the right budget line. The €2,800/month net SRE (approximately €4,200 gross with employer contributions) appears as a "personnel charge" — compressible, substitutable, and first in line during the next cost optimisation programme.

◆ WHAT THE VM DOES — AND DOES NOT DO

What the VM does: it runs. It consumes. It sends a bill on the 1st of the month. It does not take holidays. It does not ask for a raise. It does not get sick. It does not resolve incidents.

What the VM does not do: it does not configure itself. It does not detect the anomaly at 3am before it becomes an incident. It does not understand why the payment service has been slow for 20 minutes. It does not write the post-mortem. It does not train the junior. It does not read the Terms of Service before signing. It does not tell the CIO that the architecture is drifting toward irreversible lock-in.

What the SRE does: everything the VM cannot do — plus the VM itself.

◆ THE COST OPTIMISATION PROGRAMME — THE CLASSIC SCENARIO

Quarter 3. Management decides to optimise costs. Two lines are on the table:

Line A: "Cloud Infrastructure — 47 GCP instances — €142,000/month." → The CIO says: "We cannot touch that, it is our entire infrastructure."

Line B: "SRE Team — 4 people — €16,800/month gross with contributions." → The CHRO says: "We could reduce to 3 people and outsource the rest."

The €3,000/month VM is untouchable. The €2,800/month net SRE is optimisable. The machine is protected by its complexity. The human is exposed by their visibility.

HUMAN

◆ THE SARCASTIC FORMULA

The organisation rents a machine for €3,000/month that it will never own, under California jurisdiction, with non-cancellable commits. It pays the human who keeps it running €200 less per month. And when it comes time to choose between the two, it keeps the machine.

This is creative accounting. The machine is in the infrastructure budget — protected. The human is in the HR budget — exposed. Even though removing the human will cost ten times more in unresolved incidents, in training their replacement, and in accumulated technical debt.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE ETERNAL RENT — INVERTED DEPRECIATION

AN ASSET THAT APPRECIATES FOR THE VENDOR · A CASH FLOW THAT NEVER STOPS FOR YOU

A normal asset depreciates over time. A car bought for €20,000 is worth €8,000 five years later. A server bought for €15,000 is worth €3,000 five years later. Value decreases — that is the nature of physical assets. The hyperscaler VM has no residual value because it is not an asset — it is a service. And a service does not depreciate. It appreciates — because the vendor can revise prices "at any time".

◆ PROPERTY VS CLOUD — THE ANALOGY THAT HURTS

Imagine a landlord who says: "I am renting you this flat for €3,000/month. You will never be able to buy it. I can revise the rent whenever I wish, with no guaranteed notice period. If you want to move out, you pay transfer fees for every piece of furniture you take. The contract is subject to California law. And if the US justice system wants to visit without telling you — the law sometimes prevents me from informing you, but not always."

No CFO in the world would accept these conditions for a flat. Every CFO in the world accepts them for cloud infrastructure. Because it is called "OpEx" and "flexibility" instead of what it actually is: an eternal rent with unfavourable contractual conditions.

◆ THE REFURBISHED GRADE A — THE FINAL ARGUMENT

A Dell PowerEdge R750 Refurbished Grade A — same specs, 3 years younger, tested and certified — costs €4,000 to €6,000. Depreciated over 3 years: €140 to €170/month. With opex (power, network, colocation): approximately €270/month. Over 3 years: €9,720.

The GCP VM over the same period: €108,000.

Ratio: 11×. For that price, you could hire a full-time senior SRE to manage your Refurbished infrastructure — and you would still have €87,000 of annual budget left for something else.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · WHAT THE BOARD DOES NOT SEE

VM IN THE INFRA BUDGET · SRE IN THE HR BUDGET · BOTH DO THE SAME THING

◆ CLOUD'S CREATIVE ACCOUNTING

Cloud pulled off a remarkable accounting sleight of hand: it transformed a CapEx expense (server purchase) into an OpEx expense (monthly subscription). Finance directors love OpEx — it smooths spending, simplifies accounting, avoids balance sheet capitalisation. What finance directors did not see is that cloud is the most expensive OpEx in the history of computing.

A server purchased as CapEx appears on the balance sheet as an asset, then depreciates — reducing the tax base each year. A VM as OpEx is simply an expense — with no asset counterpart, no depreciation, no residual value. It is the OpEx that enriches the vendor and impoverishes the client, without the client ever accumulating anything.

◆ TOTAL COST OF OWNERSHIP — THE REAL COMPARISON

When an organisation compares "cloud vs on-premise", it often compares the visible cost of cloud (the monthly bill) to the total cost of on-premise (purchase + staff + maintenance + datacentre). It forgets that cloud also has a staff cost — AWS-certified DevOps cost more than classic system administrators. It forgets egress fees, Premium support, migration consulting fees. And it always compares an already-operational cloud to a hypothetical on-premise yet to be built — never to an already-built on-premise.

◆ THE QUESTION THE BOARD SHOULD ASK

"If I took the €3,000/month I pay for this GCP VM and invested it in a bare-metal server, what would happen?"

Months 1 to 5: buy the server (€15,000 — 5 instalments of €3,000). Same monthly budget.

Month 6: the server is paid off. Monthly cost drops to €150. Monthly saving: €2,850.

Over 5 years: total saving of €156,000. With this money, you can recruit an additional senior SRE, buy two Refurbished backup servers, fund three years of training for your team, and still have enough left for the coffee machine.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE PROPOSAL — REFURBISHED + SENIOR SRE

SOVEREIGN INFRASTRUCTURE FOR THE PRICE OF ONE AND A HALF VMS

◆ WHAT YOU CAN HAVE FOR €3,000/MONTH — THE REAL COMPARISON

For €3,000/month at GCP: an n2-highmem-32 VM. You do not own it. It runs under California jurisdiction. When you want to leave, you pay egress fees. In 5 years you will have spent €180,000 and you will have nothing.

For €3,000/month for 5 months + €150/month thereafter: a Dell PowerEdge R750 bare-metal server you own, under your jurisdiction, in your colocation or with your sovereign managed service provider. After 5 months, your monthly cost drops to €150. In 5 years you will have spent €24,000 and you have an asset.

For €5,000/month (VM + €2,000 extra): the Refurbished Grade A server + a full-time senior SRE who manages it, documents it, trains your team and sleeps better at night. In 5 years you have sovereign infrastructure, a competent team, and an SRE who knows your infrastructure from the inside.

◆ THE REFURBISHED SRE — THE UNCOMFORTABLE ANALOGY

If Refurbished Grade A logic were applied to hiring: recruit an SRE with 10 years of experience (Refurbished — has already resolved hundreds of production incidents) rather than an AWS-certified junior (new, brilliant but without lived experience). The senior SRE costs €3,500 net instead of €2,800. Delta: €700/month.

What the senior SRE brings extra: they prevent on average 2 to 3 incidents per year that would each have cost 4 to 8 hours of downtime × business impact. They optimise the infrastructure to reduce costs. They document what they do. They train juniors. They sleep better at night — and so do you.

ROI of the €700/month delta: incalculable. But certainly positive.

HUMAN

The €3,000 VM does not resolve the 3am incident.
The €2,800 net SRE does.
And yet it is the VM that is in the protected budget.
The next time your CIO says "we cannot touch the cloud infrastructure",
ask them how much the SRE who prevents it from going down costs.
Then watch them calculate.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE CLOUD CORPORATE GLOSSARY

WHAT THEY SAY · WHAT THEY MEAN

THEY SAY	THEY MEAN
"OpEx vs CapEx — the flexibility"	Eternal rent with no purchase option, revisable "at any time" (GCP §2.6). What the CFO would not accept for their HQ, they accept for servers.
"Cloud-native"	Incompatible with anything that is not this cloud. "Cloud-native AWS" is the polite word for "AWS lock-in". They call it a skill. The corpus calls it a dependency.
"Infinite scalability"	You pay infinitely. The scalability is real. So is the invoice. And when you scale down, noncancellable commits block you.
"Multi-AZ resilience"	You pay twice for something nobody really tests. See us-east-1, December 2021. Note: the centralised control plane that orchestrates failover was itself affected.
"Our cloud journey"	Migration toward dependency, presented in the boardroom with blue slides and upward arrows. Duration: infinite. Budget: growing. Return on investment: unmeasured.
"Continuous innovation"	New proprietary services creating new dependencies. Each "innovation" is a new dialect your code will have to speak. When AWS deprecates — you migrate, at your expense.
"Serverless"	Server-elsewhere-and-metered-to-the-nanosecond. The servers exist. They belong to someone else. You cannot see them — you could not see them from us-east-1 either.
"Zero-ops"	Zero internal skills, total dependency. Terminal stage of the amputation documented in the corpus. When it is time to leave, there will be nobody left to lead the migration.
"Best of breed"	The best within their catalogue, among their services. When all the "best tools" are AWS/GCP/Azure, "best of breed" means "maximum lock-in".
"We're sorry to hear you're experiencing difficulties"	Your ticket is in the queue. A bot replied. A human will read it in 48h — or 8h if you pay the Enterprise plan at €50,000/year. Your managed service provider picks up the phone.

HUMAIN

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE FICTIONAL BOARD MEETING — SIMULTANEOUS TRANSLATION

WHAT THE CIO SAYS · WHAT IT MEANS · WHAT THE CFO SHOULD HEAR

Scene: meeting room. Tuesday 2pm. Infrastructure roadmap presentation. The CIO has slides. The figures have been rounded.

THE CIO SAYS	TRANSLATION
"Our cloud journey has enabled unprecedented agility through a highly available cloud-native multi-AZ architecture."	We migrated to AWS 3 years ago, we pay €142,000/month, the code speaks DynamoDB, and the multi-AZ failover test has not been done in 14 months.
"We have optimised our OpEx through a targeted Reserved Instances strategy on our critical workloads."	We signed noncancellable 3-year commits. If we want to leave, we still pay. We call it "optimisation" because it is 20% cheaper than on-demand — but still 7.5× more expensive than bare-metal.
"Our multi-region BCP guarantees enterprise-grade resilience with a documented 15-minute RTO."	The BCP is documented. It has not been production-tested in 8 months. The secondary region Terraform has diverged. The secondary load balancer SSL certificate expires in 3 weeks. Estimated real RTO: 4 hours.
"We have a team of 4 highly skilled SREs ensuring the availability of our serverless infrastructure."	4 SREs at €2,800 net/month manage infrastructure costing €142,000/month. The machines cost more than the team that keeps them running. And in the Q3 optimisation plan, it is the team that is "optimisable".
"Our strategic partnership with AWS ideally positions us to accelerate our digital transformation."	We signed an Enterprise Agreement. An "AWS Account Manager" calls us every month to propose new services. We call it a partnership. AWS calls it a captive client. The Terms of Service call it Section 14.12 — California jurisdiction.

HUMAN

The CFO listened. They nodded. They did not do the calculation.
The Opération Dindon corpus did it for them.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · JUNE 2026

THE ECONOMY OF SILENCE

Uptime as the Only Metric

From the permanent fire to the infrastructure that works on its own

RATIO

◆ **FOREWORD — A REQUEST FOR READING**

This document advances a thesis that is simple and, on the surface, provocative: the best SRE is often the one who appears to do nothing, because their infrastructure works for them rather than the other way around. This is not a celebration of laziness. It is a thesis about metrics: an SRE's value is not measured in visibly busy hours, but in the real availability of the service they built. The reader is invited to judge this thesis on its internal coherence, not on the reaction it might trigger at first glance.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine Operation Dindon

TABLE OF CONTENTS

◆ Section 1 · Uptime as a Value Metric, Not Visible Time
◆ Section 2 · The Absence of Panic as Proof of Competence
◆ Section 3 · Cognitive Recovery as a Condition for Performance
◆ Section 4 · Timesheets, From Bad Metric to Competitive Climate
◆ Section 5 · Conclusion — What the SRE Actually Owes the Organisation

◆ READING METHOD

Like the rest of this corpus, this document systematically distinguishes structural cause from visible symptom, and relies on concepts documented in professional SRE literature rather than personal anecdote. The humour used in Section 3 takes nothing away from the rigour of the argument: it serves to name a real and widely shared practice rather than to hide it behind sanitised managerial vocabulary.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · THE VALUE METRIC

UPTIME AS A VALUE METRIC, NOT VISIBLE TIME

A sysadmin who appears constantly busy, who fights incidents on repeat, who responds to alerts at all hours, gives an impression of dedication and high workload. This impression is misleading. A well-designed system does not produce serial incidents: it anticipates failures before they become visible. Visible activity is therefore not proof of good work. Most often, it is the symptom of an earlier failure to anticipate.

◆ THE CAUSE / SYMPTOM DISTINCTION

This point requires precision so as not to hurt anyone: no professional is, alone and in the moment, responsible for all the technical debt they fight against. The SRE who spends their days firefighting often inherits infrastructure built before them, under deadline pressure, without sufficient engineering budget. The observation is therefore never a criticism of the person in post. It is a structural observation: when visible activity is constant, something upstream — architecture, budget, prioritisation — did not receive the attention it needed at the right time.

Site Reliability Engineering, as documented by the teams who formalised the discipline at Google, names this phenomenon precisely under the term "toil": manual, repetitive operational work that grows proportionally with system scale and produces no durable value once executed. The SRE discipline sets an explicit limit on this work — it must not exceed a reasonable share of total time, the rest being devoted to engineering that reduces this repetitive work in the future.

◆ NASSIHA — WHAT THIS THESIS DOES NOT SAY

This thesis does not say that fighting incidents is a sign of personal incompetence. It says that a system producing continuous incidents signals a structural underinvestment in prevention, and that measuring an SRE's value by their visible level of busyness accidentally rewards the symptom rather than the cause.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE ABSENCE OF PANIC

THE ABSENCE OF PANIC AS PROOF OF COMPETENCE

It is midnight on a Black Friday. A company's e-commerce traffic has multiplied eightfold within hours. In the virtual war room, part of the technical team cycles through manual restarts, watches dashboards blinking red, and negotiates live with the cloud provider's support to secure more capacity. The SRE responsible for the payment infrastructure is not in that room. They are asleep.

They are not asleep out of negligence. They are asleep because the previous six months were dedicated to exactly this moment: load tests simulating ten times normal traffic, automated failover drills repeated in pre-production, an auto-scaling policy validated against real scenarios rather than theoretical assumptions. The Black Friday spike is not an unforeseen event for their system — it is a scenario already rehearsed several times, under conditions harsher than reality.

> **◆ THE INVISIBLE WORK BEHIND THE SILENCE**
>
> The SRE's silence during the event is not the absence of work. It is the visible conclusion of invisible work completed months earlier — capacity planning, chaos engineering, automated failover. The difference between the panicking team and the sleeping SRE is not a difference in individual competence in the moment. It is a difference in when the effort was invested: before the incident rather than during it.

This scenario, deliberately composed to illustrate the principle rather than to recount one specific real case, reflects a documented and widely taught practice in reliability engineering: regular failure drills, conducted upstream and outside periods of criticality, replace improvised crisis management with already-rehearsed execution. Silence in production is not a mystery. It is the measurable result of work correctly sequenced over time.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · COGNITIVE RECOVERY

COGNITIVE RECOVERY AS A CONDITION FOR PERFORMANCE

Let's say out loud what many have practised quietly since remote work became the norm. The video call where one wears a shirt, sometimes a tie, above the belt, and simple underwear or shorts below the camera frame, has become a universal scene of remote work — funny because it is true, and true for a reason that goes beyond mere convenience.

> **◆ WHAT THIS SCENE REVEALS**
>
> This scene is not a lapse in professional discipline. It is the expression of a real need: body and mind need moments of detachment to remain performant over time. One does not resolve a complex incident or write reliable code under continuous mental pressure. Micro-recovery — a short nap, a screen-free break, a moment of disconnection — is not time stolen from the employer. It is the physiological condition for the useful work time that follows.

Sleep and cognition science have documented this mechanism for decades: a short nap improves alertness, working memory and problem-solving ability in the hours that follow, unlike forced sustained attention, which progressively degrades the quality of judgment. An SRE who takes twenty minutes of sleep in the afternoon, or a clean break between two tasks, does not work less over the day. They avoid the degraded mode of operation that produces configuration errors, rushed on-call decisions, and ultimately new incidents.

> **◆ NASSIHA — THE LIMIT OF THIS PERMISSION**
>
> This section does not argue for an absence of availability. An SRE remains responsible for responding to a genuine incident, both during working hours and on call. What is defended here is the recognition that, between two real activity peaks, mental detachment — in whatever form it takes — is not a failing, but a performance condition documented by the scientific literature on cognition.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

TIMESHEETS, FROM BAD METRIC TO COMPETITIVE CLIMATE

The activity report, in its most common form, asks people to justify time spent rather than results produced. This is, first and foremost, a metric error. Labour sociology documents this phenomenon under the name presenteeism: valuing visibly occupied time over actually produced results, which pushes individuals to favour the appearance of activity over effectiveness.

◆ THE FIRST STAGE — THE METRIC ERROR

For an SRE whose infrastructure runs incident-free for an entire week, filling in a classic timesheet becomes a contradictory exercise: either they invent activity to justify their 35 hours, or they honestly document that everything worked without them, which in a culture of visible time can be read as a negative signal rather than as proof of successful upstream engineering work.

This first imbalance, once embedded in a team's culture, does not remain a mere administrative curiosity. It produces a second, deeper effect: a competitive climate around appearing busy. If the recognised metric rewards visible activity, team members are structurally incentivised to display themselves as more loaded than their neighbour, rather than to cooperate so that the shared infrastructure has as few incidents as possible.

◆ THE SECOND STAGE — THE CLIMATE THAT FEEDS ITSELF

This competitive climate becomes a problem in its own right, independent of the metric that produced it. Once "appearing busy" becomes an individual visibility strategy, cooperation — sharing an automation that would reduce everyone's workload, documenting a solution so no one has to rediscover it — loses perceived value against simply demonstrating personal activity. The system meant to measure contribution ends up discouraging the most useful contribution of all: the one that makes future work unnecessary.

The timesheet is therefore not merely a symptom of an initial bad metric. Once this climate is established, it becomes a structural problem in its own right, one that keeps producing its effects even if one attempts to fix its wording without changing the culture that gave rise to it.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · CONCLUSION

WHAT THE SRE ACTUALLY OWES THE ORGANISATION

An SRE does not owe their organisation 35 hours of visible presence per week. They owe the optimal availability of the service they are responsible for. These two obligations rarely overlap in practice, and this document has tried to show why: a well-designed system keeps its guardian silent, while a poorly designed system keeps its guardian visible, busy, and apparently indispensable at every moment — for the wrong reasons.

◆ SYNTHESISING THE FOUR MECHANISMS

The absence of incidents signals successful anticipation, not an absence of work. Silence during a load spike is the visible consequence of invisible effort completed months earlier. Cognitive recovery, in whatever form it takes, is a condition for performance rather than stolen time. And the visible-time metric, once established, does not remain a simple measurement error: it becomes a climate that discourages the most useful long-term cooperation.

The practical consequence of this analysis is not the abandonment of all SRE performance measurement. It is a shift in criteria: measuring the service's actual availability, the reduction of repetitive work over time, and the quality of anticipation — rather than the volume of visibly occupied hours. An SRE who appears to do nothing because their infrastructure works for them is not an employee falling short. They are the most tangible proof that their earlier work was done correctly.

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE TICKET AND THE TALENT

Infrastructure roles nomenclature

SysOps · NetOps · OpInfra

◆ **CONTEXT OF THE STUDY**

This study proposes a structured nomenclature for digital infrastructure roles in three categories — SysOps, NetOps, OpInfra — and six hierarchical levels. It directly responds to "The Infrational Crisis" by placing precise words on realities that DevOps vocabulary has progressively erased. The founding rule: never go beyond this naming, at the risk of losing meaning and producing the dissolution the corpus has documented.

Amine RAITI — Infrastructure Architect & SRE

Teaching since 2006 · Electricity to Kubernetes · All audiences

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · WHY A NOMENCLATURE — AND WHY NOW

WORDS BEFORE STRUCTURES

A roles nomenclature is not a bureaucratic exercise. It is a political act — in the etymological sense: an act that concerns the city, the community, the organisation of professional life together. Naming infrastructure roles correctly is laying the conditions for their recognition, remuneration, protection and transmission.

◆ WHAT THE ABSENCE OF NOMENCLATURE PRODUCES — THE CORPUS DIAGNOSIS

"The Infrational Crisis" documented how word inflation — DevOps, Infrastructure as Code, XxxOps — emptied roles of their substance. "The Infrational Loop" documented how this title dissolution produces salary compression that pushes seniors out. "The State and the Invisible Body" documented how the absence of a recognised nomenclature deprives the infrastructure body of any institutional protection. These three pathologies share a common cause: nobody placed precise words on precise realities, and left the market to name at its convenience.

◆ WHAT THE RNCP DOES — AND WHAT IT DOES NOT DO

The RNCP already partially recognises infrastructure roles. RNCP37682 (Systems and Networks Senior Technician, level 5) and RNCP35594 (Systems, Networks and Databases Administrator) exist and are registered. The problem is twofold: these certifications systematically fuse systems and networks into a single title — erasing the functional distinction between the two domains — and they do not cover higher levels (engineer, lead, architect) with the same precision. This corpus's proposal completes the RNCP by offering a coherent taxonomy across the full pyramid.

◆ NASSIHA — THE NOMENCLATURE IS NOT AN END IN ITSELF

A nomenclature without legal force is a catalogue. It has value only if adopted in collective agreements, public procurement frameworks, CIO job descriptions and training programmes. This study places the words. "The State and the Invisible Body" placed the mechanisms to give them force. Both are necessary.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE THREE CATEGORIES

SYSOPS · NETOPS · OPINFRA

The nomenclature is organised into three functional categories. They are not silos — they are primary competence domains that define where the engineer's centre of gravity lies. A profile may operate across two categories simultaneously — in small organisations, this is often necessary. But their title remains anchored in their primary category.

◆ SYSOPS — SYSTEMS OPERATIONS

Scope: everything touching operating systems, physical and virtual servers, low-level application services, storage, virtualisation, backup and restoration. It is the layer that transforms metal into usable service.

What SysOps is not: SysOps is not application development. A SysOps engineer may write automation scripts — but does not develop business applications. The boundary is the service layer: below it is SysOps. Above it is the developer.

Physical anchor: a SysOps engineer has touched a physical server. They know what RAID, a storage bay, and a hypervisor are. They can diagnose a hardware fault before opening a console.

◆ NETOPS — NETWORK OPERATIONS

Scope: everything touching networks — cabling, switching, routing, firewalls, VPN, WiFi, WAN, BGP, MPLS, DNS, DHCP, network monitoring. It is the layer that connects systems to each other and to the outside world.

What NetOps is not: NetOps is not security in the broad sense — SecOps is a distinct specialisation that builds on NetOps but is not a sub-category of it. A NetOps engineer secures the network. They are not the CISO.

Physical anchor: a NetOps engineer has cabled a switch. They can read a rack diagram, trace a backbone cable, diagnose packet loss on copper before blaming software.

◆ OPINFRA — CROSS-DOMAIN INFRASTRUCTURE OPERATIONS

Scope: the cross-domain layer that operates infrastructure as a whole — SREs, infrastructure architects, senior operations engineers who master both domains and operate at a level of responsibility over overall system availability.

What OpInfra is not: OpInfra is not DevOps. An OpInfra profile may automate, may write infrastructure code — but their centre of gravity is operational reliability, not feature delivery.

Physical anchor: an OpInfra engineer has documented mastery of the physical layer in at least one domain (SysOps or NetOps) before accessing engineer level.

Amine RAITI · CC BY-NC-SA 4.0

FROM SUPPORT TO ARCHITECTURE — THE COMPLETE PYRAMID

◆ THE PYRAMID — COMPLETE TABLE

LEVEL 1 — SUPPORT / PROXIMITY

Support Technician · Field Technician

First user contact. First-level diagnosis. Escalation to level 2. Does not touch server or network configurations in production without supervision. Maps to RNCP37682 block 1 partial.

LEVEL 2 — DOMAIN TECHNICIAN

Systems Technician (SysOps) · Network Technician (NetOps)

Daily operations. Maintaining operational condition. Operates on systems or networks under instructions and procedures. Contributes to documentation. Maps to full RNCP37682 (level 5 — Bac+2 equivalent).

LEVEL 3 — DOMAIN ADMINISTRATOR

Systems Administrator (SysOps) · Network Administrator (NetOps)

Autonomous administration. Configuration design. Responsibility for infrastructure sub-systems. Level 2 incident management. Maps to RNCP35594 (level 6 — Bac+3/4 equivalent).

LEVEL 4 — DOMAIN ENGINEER

Systems Engineer (SysOps) · Network Engineer (NetOps) · SRE Engineer (OpInfra)

Intermediate architecture design. Responsibility for reliability of a defined perimeter. Supervision of lower levels. Writing procedures and runbooks. RNCP level 6 to 7 (Bac+3 to Bac+5 equivalent).

LEVEL 5 — LEAD / PRINCIPAL

Systems Lead · Network Lead · Principal SRE

Senior technical reference. Complex architecture design. Team mentoring. Structural technical decisions. Not necessarily a line manager — technical expertise first.

LEVEL 6 — ARCHITECT

Systems Architect · Network Architect · Infrastructure Architect (OpInfra)

Cross-domain vision. Infrastructure design as a whole. Technical governance. Interlocutor of the CIO and leadership. Responsibility for coherence and medium-term evolution.

RATIO

◆ NASSIHA — THE LEAD IS NOT THE MANAGER

The Lead / Principal level is a technical — not managerial — progression path. A Principal SRE who does not wish to manage teams must not be forced to become a manager to access the remuneration level corresponding to their competence. Dissociating technical progression from managerial progression is a condition for retaining the best technical profiles.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE MULTIPLE HATS RULE

ONE PROFILE CAN HOLD SEVERAL FUNCTIONS — NOT SEVERAL TITLES

The multiple hats rule is the pragmatic response to the reality of small and medium organisations. It does not contradict the nomenclature — it applies it intelligently in contexts where resources do not allow strict domain specialisation.

◆ THE PRINCIPLE — MULTIPLE FUNCTIONS, ONE TITLE

A profile may exercise several nomenclature functions simultaneously. A Systems Administrator in a small town hall may also manage the local network, provide first-level support, and administer workstations. This is real and legitimate. But their title remains "Systems Administrator" — not "IT Manager", not "Digital Project Coordinator", not "Infrastructure Lead". These vague titles have no content in the nomenclature and open the door to dilution.
Functions exercised outside the primary title are documented in the job description as "complementary functions exercised according to organisational needs" — not as an additional title.

◆ CONCRETE EXAMPLES BY ORGANISATION SIZE

Small town hall (5,000 inhabitants): 1 profile, Systems Administrator (SysOps, level 3). Complementary functions exercised: local network management, user support, workstation administration. Not called "IT Manager" — called Systems Administrator with extended responsibilities.

SME (50 employees): 1 Systems Administrator + 1 Support Technician. The first manages infrastructure and networking. The second manages user support and workstations. Two people, two titles, no confusion.

Mid-size company (500 employees): 1 SRE Engineer (OpInfra) + 2 Systems Administrators + 1 Network Administrator + 2 Support Technicians. Specialisation becomes possible and necessary. Domain boundaries are respected.

Large IT department (5,000+ employees): full pyramid by domain. Leads and Architects emerge. OpInfra transversality is embodied by a dedicated SRE team.

◆ WHAT THE RULE EXPLICITLY FORBIDS

— Calling a profile "DevOps engineer" without documented physical layer competence
— Calling a profile "cloud architect" without having designed network or systems architecture
— Creating composite titles not defined in the nomenclature ("Infrastructure & Cloud Expert", "Ops Lead", "Tech Lead DevSecOps")
— Using a higher-level title to save on recruiting a properly-levelled profile

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE BOUNDARIES THAT PROTECT MEANING

WHAT THE NOMENCLATURE DOES NOT NAME — AND WHY

A nomenclature that says what it names must also say what it does not name. Boundaries are as important as definitions — they protect meaning against progressive dilution.

◆ DEVOPS IS NOT A TITLE IN THIS NOMENCLATURE

"DevOps engineer" does not exist in this nomenclature. Not because the DevOps philosophy is rejected — collaboration between development and operations is a real value. But because "DevOps" is not a role: it is a work culture, an approach, a philosophy. One does not name a role after a philosophy. One names a role after what the person concretely does — the systems they administer, the networks they configure, the infrastructure they operate. A "DevOps engineer" who administers Linux servers is a Systems Engineer (SysOps) working in a DevOps culture. The title says what they do. The culture is documented in team practices, not in the title.

◆ SECOPS, FINOPS, MLOPS — ARE NOT CATEGORIES

SecOps is a specialisation — a Systems or Network Engineer with a security specialisation. They were called "Systems Engineer, security specialisation" or "Network Engineer, security specialisation". FinOps is a cloud financial governance practice — not an infrastructure role. MLOps is a hybrid developer/operations profile oriented toward machine learning — belonging more to the developer category than the infrastructure category. These specialisations exist. They do not justify creating a new category in the infrastructure nomenclature.

◆ CLOUD ENGINEER IS NOT A TITLE IN THIS NOMENCLATURE

"Cloud Engineer" or "Cloud Architect" are marketing titles that describe an execution context (the cloud), not an operational competence. A Systems Engineer who works primarily on AWS is a Systems Engineer — not a "Cloud Engineer". The precision of the competence domain (Linux, Windows, virtualisation, containerisation) and hierarchical level (engineer, lead, architect) says everything needed about the profile. The word "Cloud" in the title says where they work — not what they know how to do.

RATIO

◆ THE FOUNDING RULE — DO NOT GO BEYOND THIS NAMING

Any organisation adopting this nomenclature commits to not creating titles beyond those defined. No "Super Senior Cloud DevOps Infrastructure Engineer". No "SRE Platform Tech Lead". The existing words in the nomenclature suffice to describe any real infrastructure profile. If a title cannot be expressed in the nomenclature's words, either the profile is poorly defined, or the title was invented to mask a classification problem.

Amine RAITI · CC BY-NC-SA 4.0

FROM WORD TO ACT — THREE ADOPTION VECTORS

A nomenclature has value only if adopted. Adoption passes through three distinct vectors that reinforce each other. Each can act independently — together they create critical mass.

◆ VECTOR 1 — ORGANISATIONS THEMSELVES

Any IT department, digital services firm, local authority or company can adopt this nomenclature tomorrow, without waiting for legislation or collective agreements. It replaces its vague titles with nomenclature titles in its job descriptions, org charts and job postings. This unilateral decision has three immediate effects: it clarifies responsibilities internally, it improves recruitment precision, and it contributes to building the critical mass that will make the nomenclature visible on the market.

◆ VECTOR 2 — COLLECTIVE AGREEMENTS AND RNCP

The nomenclature can be proposed to digital sector social partners as a basis for revising classification grids. Syntec, which covers the majority of IT services firms, has a classification grid that does not clearly distinguish SysOps / NetOps domains. A revision integrating the nomenclature would produce salary scales by level and domain — mechanically reducing the compression documented in "The Infrational Loop". Simultaneously, the titles can be proposed to France Compétences for RNCP registration — complementing existing certifications, not replacing them.

◆ VECTOR 3 — TRAINING ORGANISATIONS

BTS SIO, DUT Networks & Telecoms and Bac+3/4 infrastructure training programmes can align their career outcome labels with the nomenclature. Instead of promising a "systems and networks technician or DevOps engineer position", training promises a "Systems Technician position (SysOps, level 2)" with a documented path toward "Systems Administrator (level 3) then Systems Engineer (level 4)". Pathway readability is the condition of filière attractiveness — particularly for female audiences documented in "The Invisible Amputation".

◆◆◆

Words are the least costly and most structuring acts. Naming a role correctly is giving it back its value. This nomenclature is an act of resistance against dissolution — and an act of construction for those who come after.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE NEWSPEAK THAT COSTS DEAR

Anatomy of a $580 Billion Vocabulary

How Ten Words Destroyed Teams, Budgets and Sovereignties

RATIO

◆ **THE THESIS**

Cloud newspeak is not a communication accident. It is a coherent system that has produced measurable effects on infrastructures, teams and budgets. Each term has a precise function in the cognitive capture chain. Each term has left traces in balance sheets, org charts and Terms of Service. This study documents these effects term by term — with the figures that make them incontestable.

CLOUD MARKET	JOBS LOST	TRANSFORMATION FAIL	WATERMARK
$580Bn	**-34% Linux**	**70%**	**RATIO**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · NEWSPEAK AS AN ACCOUNTING WEAPON

TEN WORDS THAT SHIFTED BILLIONS FROM CAPEX TO OPEX WITHOUT ANYONE NOTICING

The global cloud market represented $230 billion in 2019. It represents $580 billion in 2023. This 152% growth in 4 years is not value creation — it is largely the migration of CapEx spending into OpEx spending, made acceptable by carefully chosen vocabulary.

◆ "OPEX" — THE WORD THAT MADE ACCUMULATION INVISIBLE

A €180,000 investment over 5 years goes to the board for approval. A €3,000 monthly expense flies under the radar. It is exactly the same thing financially — but not politically. Cloud transformed investment decisions into recurring subscriptions. The result: organisations spending 10 times more than they would in CapEx, never building an asset, without the figure ever being aggregated and presented as a strategic decision.

In France, the share of large enterprise IT budgets allocated to cloud rose from 18% in 2019 to 41% in 2023. The budget allocated to internal technical teams fell by 22% over the same period. The "OpEx" newspeak made these two trends invisible to each other.

◆ "PAY-AS-YOU-GO" — THE WORD THAT NORMALISED OWNING NOTHING

"You only pay for what you consume" is true in the per-second billing sense. It is false in the economic sense: you consume forever, you never build an asset, and the accumulated value of your payments enriches the vendor without ever belonging to you. Over 5 years, "pay-as-you-go" produces €180,000 of spending on a GCP instance — and €0 of asset. The same budget in CapEx produces a server, internal competence, and residual value.

◆ "SCALABILITY" — THE WORD THAT JUSTIFIED ELIMINATING ENGINEERS

"The platform scales automatically" made the question "who manages capacity?" superfluous. The implicit answer was "nobody — the platform does it". In reality, automatic upward scaling is real. Cost reduction downward is blocked by noncancellable commits. And the platform that scales automatically still needs engineers to configure it, monitor it, and resolve incidents. But those engineers were eliminated because "the platform scales".

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · NEWSPEAK AS A DESKILLING TOOL

HOW TEN WORDS DESTROYED 34% OF NEUTRAL-SKILL JOBS IN 4 YEARS

◆ "CLOUD-NATIVE" — THE DEVALUATION OF NEUTRAL SKILLS

"Cloud-native" created a skills hierarchy: those that are "native" (proprietary, recent, vendor-certified) and those that are not (neutral, deep, not certifiable by a hyperscaler). A Linux administrator with 15 years of experience is "not cloud-native". A junior AWS Solutions Architect-certified for 6 months is.

In France, job postings for "Linux system administrator" fell 34% between 2019 and 2023. Postings for "Cloud Engineer AWS/GCP/Azure" rose 187% over the same period. The market applied the newspeak hierarchy. And the average salary of an AWS-certified DevOps is 23% higher than that of an equivalent-experience bare-metal SRE — the proprietary certification is worth more than the neutral skill.

◆ "DEVOPS" — THE MERGER THAT ELIMINATED OPERATIONS

"DevOps" in its cloud version merged two roles by eliminating one. The developer absorbed operations — but not bare-metal operations, cloud operations. The "cloud DevOps" knows how to deploy on EKS, configure a GitHub Actions pipeline, and monitor with CloudWatch. They do not know what is under EKS. They cannot get out alone. The "DevOps" merger did not create more complete engineers — it created more vendor-dependent engineers.

◆ "UPSKILLING" — THE TRAINING THAT FINANCES DEPENDENCY

Cloud certification spending tripled between 2020 and 2023 — massively funding proprietary AWS/GCP/Azure certifications. Each AWS certification is an AWS service (documented in the Terms of Service, subject to a separate CPA agreement). It is terminable by AWS. It is not transferable to GCP or a sovereign managed service provider. Cloud "upskilling" is not skill-building — it is public investment in private dependency on an extraterritorial actor.

◆ THE DESKILLING PARADOX

Fewer internal engineers = more cloud dependency = more cloud budget = justification for reducing internal engineers. The loop is perfect. The "cloud-native" newspeak created a self-sustaining paradox: the organisations that most reduced their internal technical teams are those that most need help to exit cloud — and who no longer have anyone to lead the migration.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · NEWSPEAK AS A SUBSTITUTE FOR CRITICAL THINKING

HOW WORDS MADE THE RIGHT QUESTIONS EMBARRASSING TO ASK

◆ "DIGITAL TRANSFORMATION" — THE IDENTITY QUESTION THAT SHORT-CIRCUITED ANALYSIS

"Digital transformation" turned an operational question (how do we improve our systems?) into an identity question (are we a modern company?). A company not "undergoing digital transformation" is backward. This identity pressure short-circuited cost-benefit analyses that should have been conducted.

McKinsey (2022): 70% of "digital transformation" projects fail to meet their objectives. The average cost of a failed digital transformation for a large European company: €45 to €120 million. These failures are not primarily technical — they are framing failures. People answered "are we modern?" rather than "does this work better?"

◆ "RESILIENCE BY DESIGN" — THE WORD THAT MADE TESTING SUPERFLUOUS

If resilience is "by design" — built into the architecture — why test it? This logic produced documented but untested BCPs. The AWS us-east-1 incident of December 2021 took down "multi-AZ" services because the control plane was centralised. The "design" and reality had diverged. Nobody saw it because nobody tested — because it was "by design".

Gartner (2023): 35% of organisations that massively migrated to cloud are considering partial repatriation for cost and resilience reasons. This figure was 8% in 2020. "Resilience by design" produced organisations that discover their resilience limits during real incidents — not during tests.

◆ "BEST OF BREED" — THE TAUTOLOGY THAT CLOSED THE DEBATE

If the "best" is defined within the AWS ecosystem, then the "best" is always AWS. "Best of breed" means choosing the best tool for each need — but when all tools are evaluated within a single vendor's catalogue, "best of breed" becomes "maximum lock-in". The question "is this the best tool among all available tools?" was replaced by "is this the best tool in this catalogue?" The newspeak narrowed the comparison perimeter without anyone noticing.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · NEWSPEAK AS A CLOUD ACT VECTOR

HOW WORDS MADE EXTRATERRITORIAL HOSTING ACCEPTABLE

◆ "GLOBAL INFRASTRUCTURE" — THE GEOGRAPHY THAT REASSURES AND DECEIVES

"Our data is in the global cloud" sounds reassuring — modern, distributed, resilient. "Our data is on servers in Virginia, Ireland and Singapore, under US jurisdiction, subject to the CLOUD Act, with a legally possible non-disclosure order if US authorities require it" sounds different. It is exactly the same thing. The "global infrastructure" newspeak dissociated the physical question (where are the servers?) from the legal question (which law applies?) to reassure on the first by avoiding the second.

◆ "DATA RESIDENCY" — THE FALSE FRIEND OF SOVEREIGNTY

"Your data stays in Europe" is true in the physical sense — the bytes are on disks in Ireland or Germany. It is false in the jurisdictional sense — GCP Terms Section 14.12 stipulates that the applicable law is that of the State of California and the competent courts are those of Santa Clara County. "Data residency" created an illusion of physical sovereignty that masks the absence of legal sovereignty. European regulators took years to articulate this distinction — during which sensitive public administrations migrated to cloud "in Europe" believing they were protected.

◆ "COMPLIANCE" — THE TERM THAT SHIFTED RESPONSIBILITY

"AWS is GDPR compliant" means AWS has put in place mechanisms to help its clients be GDPR compliant. It does not mean that using AWS is GDPR compliant — a nuance the "compliance" newspeak erased. Compliance is a property of the complete system (infrastructure + usage + contract + jurisdiction), not of the vendor alone. By designating the vendor as "compliant", the newspeak shifted verification responsibility to the vendor — and removed the client's incentive to verify for themselves.

◆ WHAT THE NEWSPEAK PRODUCED ON SOVEREIGNTY

Three words — "global", "residency", "compliance" — made it acceptable for European states, public administrations and critical companies to host sensitive data with actors subject to the CLOUD Act. Not through malice. Through vocabulary. The newspeak achieved what diplomatic negotiation had failed to do: normalise extraterritorial dependency.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE FIGURES OF DESTRUCTION

WHAT TEN WORDS PRODUCED — MEASURED AND DOCUMENTED

INDICATOR	FIGURE	SOURCE	RESPONSIBLE TERM
Global cloud market growth 2019→2023	**+152%**	Gartner 2023	"OpEx" · "Pay-as-you-go"
Job postings "Linux Sysadmin" FR 2019→2023	**−34%**	LinkedIn / APEC 2023	"Cloud-native" · "Zero-ops"
Job postings "Cloud Engineer AWS/GCP/Azure" FR 2019→2023	**+187%**	LinkedIn / APEC 2023	"Cloud-native" · "Upskilling"
"Digital transformation" failure rate vs objectives	**70%**	McKinsey 2022	"Digital transformation"
Orgs considering partial cloud repatriation	**35%**	Gartner 2023	"Cloud-first" · "All-in cloud"
IT budget share (large FR companies) allocated to cloud 2019→2023	**18% → 41%**	Numeum 2023	"Cloud-first" · "Agility"
Internal technical team budgets (large FR companies)	**−22%**	Numeum 2023	"Zero-ops" · "Managed services"
Cloud certification training spend 2020→2023	**×3**	Caisse des Dépôts 2023	"Upskilling" · "Cloud-native"
GCP VM vs bare-metal ratio over 5 years	**7.5× to 11×**	Dindon Corpus 2026	"OpEx" · "Pay-as-you-go"

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE DICTIONARY OF THE LOSS

FULL GLOSSARY — CORPORATE DEFINITION · REAL DEFINITION · DAMAGE PRODUCED

TERM	CORPORATE DEFINITION	REAL DEFINITION	DAMAGE PRODUCED
Cloud-native	*Modern architecture designed for cloud*	Incompatible with anything that is not this cloud	−34% Linux jobs · +187% Cloud Engineers · application lock-in
Zero-ops	*Full automation, team liberation*	Zero internal skills. Terminal stage of amputation	−22% internal team budgets · exit impossible without help
Serverless	*No server to manage*	Server-elsewhere-and-metered-to-the-nanosecond	Opaque billing · ignorance of physical substrate
Upskilling	*Building team competence*	Public investment in certifications terminable by their issuer	Training spend ×3 on proprietary certs · dependency reinforced
Data residency	*Your data stays in Europe*	Physically in Europe · legally in California (GCP §14.12)	Illusion of sovereignty · sensitive public bodies migrated
Compliance	*AWS is GDPR compliant*	AWS has tools to help — usage remains your responsibility	Client de-responsibilisation · false sense of compliance
Resilience by design	*Resilience built into the architecture*	Control plane often centralised. See us-east-1, Dec. 2021	Untested BCPs · real RTO >> documented RTO · 35% repatriation
Digital transformation	*Strategic company modernisation*	Identity question short-circuiting cost-benefit analysis	70% failure rate · €45-120M average cost for large accounts
Strategic partnership	*Privileged vendor relationship*	Noncancellable Enterprise Agreement · GCP §14.12	Contractual lock-in · Account Manager = salesperson, not ally
HR optimisation	*Operational team efficiency*	Eliminating the engineers who could exit the cloud	Self-sustaining loop · exit impossible · perpetual dependency

RATIO

Amine RAITI · CC BY-NC-SA 4.0

"WE REDUCE OUR CARBON FOOTPRINT" — BY RELOCATING IT TO VIRGINIA

"GreenOps" is the ESG layer of cloud newspeak. It arrived at exactly the right moment — when sustainability teams were looking for arguments for their extra-financial reports and when regulators were starting to demand accountability on digital footprint. Hyperscalers provided both simultaneously: green vocabulary and renewable energy certificates.

◆ "CLOUD IS GREENER" — WHAT THE CERTIFICATES DO NOT SAY

Hyperscalers buy massively from the REC market — Renewable Energy Certificates. A REC means that somewhere in the electrical grid, a quantity of energy equivalent to the declared consumption was produced from a renewable source. It does not mean that the servers hosting your data run on green energy. It means the vendor's carbon accounting is balanced on paper.

The AWS us-east-1 datacentre in Virginia consumes several gigawatts. Virginia's electrical grid is approximately 30% powered by coal and natural gas. RECs purchased by AWS do not change this energy mix. They comptably offset real fossil fuel consumption. This is creative carbon accounting — exactly as OpEx is creative financial accounting.

WHAT THE SUSTAINABILITY DIRECTOR VALIDATED: "By migrating to AWS, we reduce our carbon footprint through resource pooling." What they actually did: relocate energy consumption to US datacentres partially powered by fossil fuels, with virtual carbon offsets purchased on a certificate market.

◆ THE HIDDEN INEFFICIENCY — QUERIES RUNNING IDLE

A poorly designed Serverless architecture (Lambda cold-starting on each request, containers running on standby, monitoring jobs querying APIs every 30 seconds) consumes compute continuously — and therefore energy continuously. This inefficiency is invisible in the cloud console: it appears as a cost line, not as a carbon footprint. The bare-metal architect who dimensions their server once optimises it to last 5 years. A serverless architecture can consume 10 times more CPU cycles for the same workload — and the sustainability team will never know because nobody measures kWh per request.

RAITI

THEY SAY

"Cloud is greener through resource pooling. We reduce our carbon footprint."

THEY MEAN

We relocate our energy consumption. We buy RECs to balance the carbon accounting. The sustainability team ticked the box. Nobody measures kWh per request.

Amine RAITI · CC BY-NC-SA 4.0

THE CONSULTANT NOT IN THE ROOM — WHO WROTE THE SCRIPT

The fictional boardroom of Section 7 featured the CIO and CFO. It was missing a character: the external consultant whose slide deck preceded all these decisions by 18 months. They are not in the room — but their vocabulary is everywhere in the CIO's presentations.

◆ THE INTEREST LOOP — DOCUMENTED WITHOUT ACCUSING

The mechanism is structural and does not imply individual bad faith:

Step 1: the consulting firm recommends a "cloud-first" strategy based on market analyses (Gartner, IDC) it cites abundantly. It invoices the scoping study: €150,000 to €500,000 depending on organisation size.

Step 2: the hyperscaler certifies the firm as a "top-tier partner" (AWS Premier Partner, Google Cloud Partner, Microsoft Gold Partner). This certification requires a minimum number of cloud certifications in the firm's teams — thus funding the proprietary certifications documented in Section 2.

Step 3: the firm conducts the migration. It invoices integration man-days. Complexity is its ally — the more complex the architecture, the longer the migration, the more man-days accumulate.

Step 4: the migration produces gaps vs estimates. The firm invoices remediation man-days. Then training. Then optimisation (FinOps). The relationship is structurally perpetual.

What the loop produces: the firm has no financial interest in recommending bare-metal or a local managed service provider — there is no "partnership" with these actors, no certification, no complex migration man-days. Advice is structurally oriented toward complexity.

THEY SAY

"Our independent study recommends a cloud-first strategy to accelerate your digital transformation."

THEY MEAN

Our firm is an AWS Premier Partner. We will invoice the migration, remediation, training and optimisation. Complexity is our business model. Bare-metal simplicity generates no man-days.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE BALANCE SHEET ASSET · TAX SAVINGS · THE PATRIMONIAL VALUE CLOUD DOES NOT CREATE

The VM and SRE study documented the 7.5× ratio between the GCP VM and bare-metal over 5 years. It did not document the tax dimension of this difference — which makes the comparison even more unfavourable to cloud.

◆ DEPRECIATION — WHAT CLOUD CANNOT DO

A server purchased for €15,000 is entered on the company's balance sheet as a fixed asset. It is depreciated over 3 to 5 years. Each year, a depreciation allowance (€3,000 to €5,000/year) is recorded as an expense — reducing taxable income. At the French corporation tax rate of 25%, this allowance generates a tax saving of €750 to €1,250/year.

Over 5 years, a €15,000 server generates €3,750 to €6,250 in cumulative tax savings. The real net cost of the server is not €15,000 — it is €15,000 minus the tax saving = approximately €8,750 to €11,250.

The €3,000/month GCP VM generates €36,000 of deductible expenses per year — but without creating an asset. Expenses are deductible in both cases. The difference: bare-metal creates an asset that appears in equity, increases the company's book value, can be pledged as loan collateral, can be resold, and can be presented to an investor as a sign of patrimonial solidity.

◆ THE PATRIMONY CLOUD DESTROYS

A company that migrates 100% of its infrastructure to cloud carries out a silent balance sheet destruction: fixed assets disappear, replaced by recurring expenses. In management accounting, this company has become structurally lighter — but also more fragile. It no longer has productive assets of its own. If the vendor raises prices, it has no patrimonial negotiating leverage. If it seeks acquisition, the buyer will find no infrastructure on the balance sheet — only subscription contracts and noncancellable commits.

"Cloud empties the treasury without creating patrimonial value" is not a metaphor. It is an accounting reality.

RATIO

THEY SAY

"Cloud optimises our cost structure by transforming CapEx into flexible OpEx."

THEY MEAN

We destroy our productive assets on the balance sheet, lose the tax benefits of depreciation, and create a perpetual expense with no residual value. Our balance sheet gets lighter — and our fragility increases.

Amine RAITI · CC BY-NC-SA 4.0

ALERT STORM IN A BLACK BOX · THE SRE ON THE FRONT LINE WITHOUT LEVERAGE

The VM and SRE study documented that the SRE resolves the 3am incident, not the VM. It did not document what "resolving the incident" actually means when the root cause is in the hyperscaler's internal layers — to which the SRE has no access.

◆ THE US-EAST-1 INCIDENT — DECEMBER 2021 — THE BLACK BOX IN ACTION

On 7 December 2021, AWS us-east-1 suffered a major incident affecting numerous services — Kinesis, Lambda, API Gateway, AppSync, EventBridge, and others. Thousands of SRE teams worldwide simultaneously received cascading alerts on their own services. The root cause: an anomaly in AWS's internal control plane.

What the SRE could do: look at their dashboards, see everything falling simultaneously, read the AWS status page ("We are investigating"), and wait. They had access to no internal AWS logs, no control plane metrics, no AWS runbook. The black box was closed. The SRE was on the front line of organisational stress — their phone was ringing, their manager was asking for an ETA — but they had zero technical leverage on the root cause.

What this reveals: "Zero-ops" also means zero visibility on incident root causes. The cloud-native SRE is an expert in consequences — they see the effects, they cannot diagnose causes when those causes are inside the vendor's infrastructure.

◆ THE DIAGNOSTIC ASYMMETRY — THE INVISIBLE DISTRESS

A bare-metal SRE facing a 3am incident has access to everything: kernel logs, hardware metrics, network traces, system calls. They can go down to the physical level if needed. Their diagnosis can reach "the NVMe disk in slot 3 has had elevated error rates for 48h" or "server 7's network card is dropping packets on the production VLAN".

A cloud-native SRE facing the same incident has access to their own application metrics — and stops at the surface of the vendor's API. They cannot see what happens under Lambda. They cannot diagnose a DynamoDB degradation. They cannot trace a request beyond CloudWatch logs. The lower layer is opaque by design. And this opacity is presented as an advantage ("managed service — you don't have to worry about it") until it falls — and you still don't have to worry about it because you cannot.

RATIO

Cloud newspeak has produced teams that are more stressed, less sovereign, less capable of diagnosis, and more dependent on a status page they do not control. This is the measurable result of ten words applied over ten years. And the CFO still has not calculated it.

Amine RAITI · CC BY-NC-SA 4.0

TEN WORDS · ONE RESULT · $580 BILLION AND EUROPEAN SOVEREIGNTY

◆ THE COMPLETE CAUSAL CHAIN

"Digital transformation" creates the identity urgency → "Cloud-native" validates the direction → "OpEx" makes the spend invisible → "Pay-as-you-go" normalises owning nothing → "Upskilling" funds proprietary certifications → "Cloud-native" devalues neutral skills → "Zero-ops" justifies eliminating engineers → "Data residency" reassures on sovereignty → "Compliance" de-responsibilises the client → "Resilience by design" removes the incentive to test → "Strategic partnership" lock-ins the organisation → "HR optimisation" eliminates the last engineers who could exit.

Each term reinforces the next. Exit becomes harder at every step. This is not a conspiracy — it is an ecosystem. And ecosystems do not need conspiracies to be effective.

◆ WHAT THE CORPUS NAMED BEFORE THE FIGURES

Opération Dindon documented cognitive capture, the infrational loop, loss of sovereignty, amputation of technical knowledge. This study demonstrates that these phenomena are measurable: −34% Linux jobs, ×3 training spend on proprietary certifications, 70% digital transformation failure, 35% cloud repatriation being considered. Newspeak is not the sole cause — it is the vector. It made possible decisions that would have been impossible if they had been named correctly.

If "digital transformation" had been called "migration toward extraterritorial dependency",
35% of organisations would have done it anyway. Not 100%.
If "cloud-native" had been called "incompatible with anything that is not AWS",
HR teams would have kept their Linux engineers.
If "data residency" had been called "physically in Europe, legally in California",
public administrations would have asked for SecNumCloud first.
Words have a cost. This one is measured in billions.

RATIO

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026

NoOps

Autopsy of a Word That Killed a Profession

Chronology · Asymmetry · Paradox · CAPEX/OPEX · SLA · Technical Memory

HUMAN

◆ THE CENTRAL THESIS

NoOps is the only marketing term in the history of computing that explicitly announces the elimination of a profession in its name — and was adopted without resistance by the organisations it targeted. Not a technical evolution. A unilateral technical disarmament plan accepted with enthusiasm by its own victim. This study performs the autopsy: the chronology of linguistic liquidation, the asymmetry of transferred operations, the paradox of displaced complexity, the CAPEX/OPEX accounting complicity, the senior purge, the ghost IaC debt, and the SLA asymmetry that tells the truth the marketing hides.

PAGES	BLIND SPOTS	WATERMARK	GRAPHS
6	3	HUMAN DOMINANT	0

Amine Raiti — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · *Operation Dindon · June 2026*

FROM COEXISTENCE TO NOMINAL ANNIHILATION — THE LOSS OF THE NAME PRECEDES THE LOSS OF EXPERTISE

◆ THE THREE STAGES OF LIQUIDATION

DevOps (2009) — Coexistence: "Collaboration between Dev and Ops." Both professions coexist. The infrastructure engineer is present, visible, named. Their role is recognised. The DevOps movement breaks delivery silos — not competence silos.

Serverless (2014) — Material spoliation: "No more servers to manage." Servers exist — but you no longer touch them. The infrastructure engineer becomes progressively superfluous in the narrative. Their scope shrinks. They remain in the org chart but lose decision-making ground.

NoOps (2016-2018) — Nominal annihilation: "Build, we handle the rest." The profession itself is eliminated from vocabulary. "Linux Administrator" and "Infrastructure Engineer" job postings disappear from recruitment platforms. Replaced by "DevOps Engineer", "Cloud Engineer", "Cloud Architect" — with fundamentally different required skills: AWS certified, Kubernetes, CI/CD. Not Linux. Not networking. Not storage.

◆ THE COMPLETE CAUSAL LOOP — FROM JOB TITLE TO TALENT POOL COLLAPSE

Job titles change → required certifications change → universities adapt programmes to job postings → training firms adapt curricula to certifications → in five years, the "neutral infrastructure" skill pool collapses (−34% Linux positions documented in The Newspeak). The labour market was reconfigured by job titles — which were reconfigured by certifications — which were reconfigured by hyperscalers. The complete causal loop that the Algorithmic HR Sabotage (Anatomy of Digital Perdition) documents as an effect. The semantic erasure of job titles is the cause.

◆ THE SENIOR PURGE — AUTOMATED DISGUISED ECONOMIC DISMISSAL

"NoOps" provided managerial vocabulary to dismiss people. "Not agile." "Resistant to change." "Legacy profile." These formulas — never written in dismissal letters, always present in boardrooms — transformed rare competence into professional handicap. When a senior system administrator with 20 years of network experience is pushed out in 2019, they are not replaced. The technical memory leaves with them. The organisation orchestrated the programmed loss of its own technical memory. That is the deep explanation for why companies currently cannot reverse course: the people capable of configuring a network switch, reading a pcap, opening a patch panel no longer exist in the org chart.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

YOU SURRENDERED CONTROL — BUT KEPT THE RESPONSIBILITY FOR ECONOMIC DEATH

◆ NOOPS FOR THE CLIENT = MOREOPS FOR THE HYPERSCALER

Operations do not disappear with NoOps — they migrate. AWS, Google and Microsoft employ tens of thousands of infrastructure engineers. Their datacentres have maintenance teams, network engineers, storage specialists, physical security experts. "NoOps" for the client means "MoreOps" for the cloud vendor. This is a competence capitulation: the company transfers operational expertise for its material survival to a third party it cannot audit, cannot control, and cannot read the runbooks of. It outsourced what it no longer understands.

◆ THE SLA ASYMMETRY — THE CONTRACTUAL LIE OF NOOPS · cf. Terms Under the Microscope (8p)

NoOps sells: "We manage operations on your behalf, you are at peace." The Terms say otherwise. In the event of a major hyperscaler outage — an entire AWS region down for 12 hours, a GCP incident impacting thousands of clients simultaneously — the hyperscaler contract provides:

What the hyperscaler reimburses: cloud consumption credits equivalent to a derisory percentage of the monthly invoice (typically 10-30% for a major interruption). Sometimes a few tens of euros on an invoice of several thousands.

What the client bears 100%: the actual operating loss · contractual penalties toward their own clients · data breach and its GDPR consequences · customer trust loss · commercial failure in extreme cases.

This is the legal tour de force of NoOps: externalise the operational act while retaining the entirety of the vital risk. The company surrendered control — but kept the responsibility for economic death.

◆ THE FORMULA FOR THE LEGAL DIRECTOR

NoOps does not externalise operational risk. It externalises operational work while retaining vital risk with the client. Contractually documented in every hyperscaler's Terms. The promise of serenity is marketing. The contract says the opposite.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE PROMISE WAS LESS COMPLEXITY · THE RESULT IS MORE ABSTRACT COMPLEXITY THAT NOBODY MASTERS

◆ THE CENTRAL PARADOX — ELIMINATING LOW-LEVEL OPS CREATED HIGH-LEVEL OPS

NoOps promised: less operational complexity. Deploy your code — we handle the rest. The cloud era reality (2018-2026) produces the opposite: an explosion of abstract complexity that neither developers nor retrained former infrastructure engineers fully master.

Service meshes (Istio, Linkerd) for inter-service communication. Hyper-fragmented IAM policies with hundreds of roles, permissions and conditions. Distributed observability (traces, metrics, logs in three different tools). Real-time cost management with alerts nobody monitors. Kubernetes clusters with hundreds of deployments whose origin nobody knows. NoOps did not eliminate the need for Ops — it made Devs responsible for an invisible infrastructure they do not master, without giving them the skills to do so.

◆ GHOST IAC DEBT — THE UNAUDITED SCRIPT GRAVEYARD

Application code is designed to be modified, updated, sometimes discarded. Physical infrastructure follows a lifecycle of maintenance, depreciation and mechanical resilience. NoOps, by making infrastructure seem like an ephemeral script applied with a click, destroyed the notion of Operational Condition Maintenance (MCO).

The result: thousands of cloud resources instantiated by obsolete IaC scripts, abandoned in production because no developer dares delete them for fear of bringing down the application. The Terraform resource instantiating a Load Balancer unused for 18 months? It runs. It bills. It waits. Nobody touches it. The developer who created it has left. The script has no owner. NoOps transformed infrastructure into a graveyard of unaudited scripts — invisible technical debt that nobody accounts for because it looks like code, not infrastructure.

◆ THE EMPTY REMOTE — THE NOOPS SYNDROME IN PRODUCTION

When an outage leaves the hyperscaler-managed environment — network saturation, BGP routing issue, peering incident — nobody in the organisation knows how to diagnose. Mastery of low-level network protocols, packet analysis, machine thermal behaviour: all left with the dismissed seniors. The organisation holds a remote control. It no longer knows what is behind the screen.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · CAPEX/OPEX COMPLICITY · THE PURGE · AI AS FINAL VECTOR

THE DISASTROUS ACCOUNTING TRADE-OFF · HUMAN CAPITAL DESTROYED · ALGORITHMIC PRODUCT PLACEMENT

◆ THE ACCOUNTING FALSE FRONT — WHY NOOPS WAS ADOPTED WITHOUT RESISTANCE

NoOps aligned perfectly with the short-term accounting vision of finance teams. Internal infrastructure teams were fixed cost centres — salaries, employer contributions, training, hardware renewal: visible CAPEX on the balance sheet. Managed cloud services were variable consumption costs — OPEX on the income statement, more flexible, more "agile" in financial vocabulary.

The CFO who validated NoOps to convert CAPEX to OPEX thought they were making a rational decision. The trade-off was accountably defensible short-term. Medium-term, it is disastrous: the "savings" on infra team salaries were far exceeded by uncontrolled cloud invoices, unanticipated exit costs, and the absence of Capacity Planning that would have allowed forecasting them.

◆ THE SENIOR PURGE AS ACCOUNTING TRADE-OFF — HUMAN CAPITAL DESTROYED

The decision to eliminate infra teams was not only organisational. It was financially motivated: a senior system administrator costs €60,000 to €90,000 per year. Eliminating three senior positions = saving €200,000 in payroll. HR has the vocabulary: "legacy profile", "not agile", "resistant to change." Finance has the numbers. The board approves.

What nobody accounted for: the value of the human capital destroyed. The memory of architectures. The knowledge of system dependencies. The diagnostic capability for outages outside the managed perimeter. These assets do not appear on the balance sheet — their destruction neither. Until the outage happens.

◆ GENERATIVE AI AS FINAL VECTOR — ALGORITHMIC PRODUCT PLACEMENT

Between 2024 and 2026, NoOps found its second wind with AI code copilots. "AI writes the YAML and configures Serverless — you no longer need an infra architect." Generative AI is not intrinsically linked to hyperscalers. But Amazon Q generates AWS code by default. GitHub Copilot (Microsoft) generates Azure code in Azure context. Not AI — algorithmic product placement. AI completes the dissolution of the last traces of hardware consciousness in the application engineer: they no longer understand what they deploy, and now they no longer even write it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE RIGHT TO NAME — REHABILITATING THE ENGINEER OF THE REAL

THE WAR FOR DIGITAL SOVEREIGNTY IS WON FIRST IN THE DICTIONARY OF MANAGERIAL CONCEPTS

◆ REIMPOSING THE TERM — THE FIRST TOOL OF ORGANISATIONAL RESISTANCE

If the loss of the name on job postings preceded the loss of expertise in competence audits — then recovering the name is the first act of resistance. Before the rack. Before the migration. Before the tender.

In org charts: "Infrastructure Engineer", "Systems Engineer", "Systems Administrator", "Network Engineer", "Network Administrator" must reappear as distinct titles — not subsets of a generic DevOps role. An infrastructure engineer is not a developer who operates. A developer is not an infrastructure engineer who codes. Two professions with two temporalities and two incompatible risk types.

In budgets: reintroduce a distinct cost line for infrastructure — not buried in a global DevOps budget. What has no budget line has no existence in the boardroom.

In tenders: require neutral and verifiable skills — Linux Foundation Certified, Red Hat, CompTIA Network+ — not only hyperscaler certifications. The AWS Solutions Architect certification does not replace mastery of physical networking.

◆ THE SOVEREIGN BRICK RULE — cf. The Dissolution of the Boundary (8p)

Every infrastructure component must be able to answer "yes" to the question: "Can it run outside AWS without a rewrite?" If no — deliberate decision, documented, assumed by both parties. Not a default choice because "it was easy to deploy." This rule is the meeting point of the infrastructure engineer and the developer — it does not say "no" to cloud, it says "conscious" of cloud.

NoOps is the only marketing term that announces

the elimination of a profession in its name.

And nobody cried out.

Amine RAITI · Operation Dindon · 2026

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE INVISIBLE DEBT

Open source as uncompensated critical infrastructure

◆ CONTEXT

The Opération Dindon corpus documented dependence on hyperscalers and semiconductors. It had not addressed the layer below: dependence on unfunded open source. Linux, OpenSSL, curl, bash — global infrastructure runs on software maintained by individuals or small teams, often volunteers or underfunded. The Log4Shell vulnerability in 2021 revealed the scale of this fragility. This study provides the structural analysis.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · OPENSSL AND THE TWO PEOPLE

GLOBAL INFRASTRUCTURE RUNS ON VOLUNTEERS

In 2014, the Heartbleed vulnerability revealed a critical flaw in OpenSSL — the cryptographic library securing the vast majority of HTTPS connections worldwide. Banks, governments, hospitals, e-commerce platforms, hyperscalers — all depended on OpenSSL. The flaw had existed in the code for two years, undetected. At the time of discovery, OpenSSL was maintained by a team of two full-time people, with an annual budget of approximately $2,000 in donations.

◆ LOG4SHELL — THE 2021 CASE STUDY

Log4j is a Java logging library embedded in millions of applications — from enterprise servers to industrial systems to video game consoles. The Log4Shell vulnerability discovered in December 2021 was rated 10/10 on the CVSS criticality scale. It potentially affected hundreds of millions of systems worldwide. Log4j was maintained by a handful of volunteers, without structured commercial funding. The emergency response to the vulnerability relied on these same volunteers, forced to work without pause for several weeks on a problem whose resolution benefited billion-dollar companies that had contributed nothing to their funding.

◆ THE STRUCTURE OF THE DEPENDENCE

The global software industry is built on an inverted pyramid: billions of dollars in market capitalisation rest, at their base, on open source libraries maintained by dozens or hundreds of people without proportional funding. This pyramid is stable as long as no critical vulnerability appears in the lower layers. When one does, it reveals that the base of the pyramid was made of cardboard.

◆ NASSIHA — OPEN SOURCE IS NOT FREE

Open source is not free. It is funded by the unpaid or underpaid time of its contributors. The "free to use" model defers the cost of development and maintenance onto individuals who bear this cost out of passion, reputation, or idealism — not out of a viable economic model. This invisibility of cost is the source of the vulnerability.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE INVISIBLE DEPENDENCY CHAIN

WHAT LIES UNDER THE CODE WE WRITE

Every developer or infrastructure engineer uses open source dependencies in their daily work — often without being aware of it. A modern web application can depend on 500 to 1,000 open source libraries, each with its own dependencies, its own maintainers, its own level of funding and maintenance. This dependency chain is largely invisible to the end user — and often to the developer themselves.

◆ XKCD 2347 — "DEPENDENCY"

A comic strip published by the webcomic XKCD depicts all the world's digital infrastructure as a stack of building blocks, one of which — tiny, at the base — is labelled "a project maintained since 2003 by a person in Nebraska". This image went viral in 2021 during Log4Shell because it accurately describes reality: somewhere in the dependency chain of every critical system, there is a library maintained by someone whose name nobody knows.

◆ THE HUMAN RISK OF OPEN SOURCE DEPENDENCE

The vulnerability of unfunded open source is not only technical — it is human. A maintainer who burns out, changes jobs, falls ill, or decides to stop abandons a library that thousands of projects depend on. This library will no longer be maintained, new vulnerabilities will no longer be fixed, new language versions will no longer be supported. The "maintainer abandonment" risk is documented but rarely integrated into infrastructure risk analyses.

◆ NASSIHA — THE PROBLEM IS NOT OPEN SOURCE

Open source is one of humanity's greatest collective creations. This study does not critique the model — it critiques the funding. A world where companies massively use open source without contributing to its funding is a world that externalises its critical infrastructure onto the goodwill of uncompensated individual maintainers. This is not sustainable.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · WHO FUNDS WHAT

THE ASYMMETRY OF VALUE AND CONTRIBUTION

Hyperscalers built multi-trillion-dollar companies on open source foundations. AWS runs on Linux. GCP depends on dozens of open source libraries. Microsoft's AI runs on Python and its ecosystems. These companies have also, for the most part, created open source contribution programmes. But the disproportion between extracted value and provided contribution remains considerable.

◆ WHAT HYPERSCALERS CONTRIBUTE

Google maintains Kubernetes, Angular, TensorFlow. Microsoft acquired GitHub and contributes massively to VSCode, TypeScript, .NET open source. Meta maintains React and PyTorch. Amazon contributes to OpenSearch. These contributions are real and significant — they do not cover the full funding debt of the open source ecosystem on which these companies depend.

◆ CRITICAL PROJECTS THAT RECEIVE NOTHING

Hyperscaler contributions concentrate on projects they control or whose visibility serves their image. Critical but low-visibility projects — compression libraries, file format parsers, network protocol implementations — remain maintained by individuals without structured funding. These libraries are used by millions of projects, including those of the hyperscalers themselves. The selection of what is funded is strategic, not altruistic.

◆ "THE SILENCE OF THE ROOMS" APPLIED TO SOFTWARE

"The Silence of the Rooms" documented that women are absent from infrastructure because they are not represented there and nobody decided to change that. The same logic applies to open source funding: invisible libraries are not funded because they are invisible, and they remain invisible because nobody decides to make them visible. Invisibility perpetuates the underfunding that perpetuates invisibility.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · MODELS THAT WORK

WHAT EXISTS AND WHAT IS MISSING

Open source funding models exist. None has solved the problem at the necessary scale. Analysing them identifies what works, what fails, and what is missing to build structural funding for critical open source.

◆ MODEL 1 — FOUNDATIONS (LINUX, APACHE, PYTHON SOFTWARE FOUNDATION)

Open source foundations collect membership fees from corporate members, employ full-time developers on critical projects, and organise governance. This model works for projects visible and widely-used enough to attract sufficient paying members. It fails for critical but low-visibility projects that cannot build a sufficient member base to self-fund.

◆ MODEL 2 — DUAL LICENSING AND OPEN CORE

Companies like HashiCorp, Elastic or Redis Labs adopted models where the base code is open source and advanced features are proprietary (open core) or the code is available under a restrictive licence for commercial use (dual licensing). These models fund development but create community tensions and compatibility breaks between versions.

◆ MODEL 3 — SOVEREIGN TECH FUND (GERMANY)

Germany created the Sovereign Tech Fund in 2022 — a public fund that directly finances the maintenance of critical open source libraries identified as national digital infrastructure. This model treats open source as public infrastructure — on a par with roads or electrical networks — and funds it accordingly. It is the model most coherent with this study's thesis. It remains marginal at European scale.

◆ NASSIHA — NO SINGLE SOLUTION

None of these models alone is sufficient. The structural solution is probably a combination: public funding for critical infrastructure libraries, mandatory contributions from large user companies proportional to their usage, and strengthened foundations for governance. It is not complicated to design. It is difficult to coordinate.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · WHAT THE SRE CAN DO

KNOWING YOUR DEPENDENCY CHAIN

The SRE cannot solve the open source funding problem alone. But they can reduce their exposure to the risk of unfunded dependence — and contribute, at their level, to the viability of the ecosystem they depend on.

◆ LEVER 1 — AUDIT YOUR OPEN SOURCE DEPENDENCIES

Know the critical libraries your infrastructure depends on, their maintenance level (date of last commit, number of active maintainers, presence of a foundation or sponsor), and their vulnerability history. This audit should be as standard as a security audit — because it is security. A tool like Dependabot or Renovate automates part of this tracking.

◆ LEVER 2 — INTEGRATE "MAINTAINER ABANDONMENT" RISK INTO RISK ANALYSES

The risk that a critical maintainer abandons their library should appear in infrastructure risk analyses, alongside the risk of supplier bankruptcy or OS end-of-support. This risk has a response: internal fork of the library, migration to a better-maintained alternative, or contribution to the maintainer's funding.

◆ LEVER 3 — CONTRIBUTE AND ENCOURAGE CONTRIBUTION

Companies that use open source can contribute through code (bug fixes, new features), through funding (donations to maintainers, foundation membership), or through time (engineers dedicated to maintaining critical libraries). These contributions are not philanthropy — they are risk management. A better-funded library is a less risky dependency.

◆ NASSIHA — CONTRIBUTION IS NOT ACCESSIBLE TO ALL

Contributing to open source requires time and skills that not all SREs and companies have. The systemic response — public funding, legal obligations of proportional contribution based on usage — is irreplaceable. Individual levers are complements, not substitutes.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

TREATING OPEN SOURCE AS PUBLIC INFRASTRUCTURE

The proposal of this study is to treat critical open source — the libraries on which the world's digital infrastructure depends — as public infrastructure on a par with roads, electrical networks or water systems. This framing is not a metaphor. It is an operational reality: if curl stops being maintained, entire sections of the world's digital infrastructure progressively cease to function. If OpenSSL is not patched, billions of secure connections become vulnerable.

◆ MEASURE 1 — CREATE A EUROPEAN CRITICAL OPEN SOURCE FUND

On the model of the German Sovereign Tech Fund, a European open source critical funding fund — endowed with a budget proportional to the size of the European digital economy — finances the maintenance of libraries identified as critical infrastructure. This fund is fed by public contributions and by mandatory contributions from companies whose turnover exceeds a defined threshold and who use these libraries in production.

◆ MEASURE 2 — MANDATORY DECLARATION OF CRITICAL DEPENDENCIES

Companies beyond a certain size should annually declare their critical open source dependencies — in the same way as declaring critical suppliers in business continuity plans. This transparency would identify at-risk libraries (few maintainers, little funding, many users) and prioritise collective funding efforts.

◆ MEASURE 3 — INTEGRATE OPEN SOURCE INTO PUBLIC PROCUREMENT CRITERIA

Every public digital services contract should integrate a criterion of contribution to critical open source in its award conditions. A provider who contributes to the funding of the open source libraries they use in their service is a more reliable provider — and their contract should reflect it.

RATIO

The invisible debt is always repaid. Sometimes in money. Sometimes in Heartbleed. Sometimes in Log4Shell. Better to choose the moment and form of repayment than to have it chosen for you.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE DEPARTURE OF THE LAST ONE WHO KNOWS

The invisible erosion of the senior in infrastructure and the knowledge that leaves with them

HUMAN

◆ CONTEXT OF THE STUDY

This study is the mirror image of "The Blue — What the Diploma Does Not Measure". The Blue is the profile who enters the sector without being recognised. The Last One Who Knows is the profile who leaves it without having been transmitted. Between the two, the same invisibility — that of tacit knowledge that formal systems cannot capture. This study addresses the erosion of the senior SRE, the progressive cognitive desertification that affects them, and what the sector loses when they leave without anyone having seen it coming.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION 1 · THE PROFILE

20 YEARS OF PRODUCTION, ZERO VISIBILITY

In every infrastructure team there is a profile that everyone knows without ever having formally named. It is the SRE or system administrator aged 45 to 55, with 20 years of seniority in the sector, who knows things no one else knows. They know why that network configuration was built that way in 2009, why that backup script has an exception on server 7, why the client A critical application cannot tolerate anyone touching the network interface between 10pm and 2am.

This knowledge is documented nowhere. It is not in the team wiki, not in the runbooks, not in the incident tickets. It lives in the head of this person, acquired incident by incident, fault by fault, sleepless night by sleepless night over 20 years. It is the most valuable knowledge the team possesses. It is also the most fragile — because it is stored in a single human point of failure.

◆ WHAT THE MARKET CANNOT VALUE

The infrastructure job market values recent certifications, fashionable technologies, the ability to talk Kubernetes or Terraform in a 45-minute interview. It values poorly — often not at all — 20 years of institutional memory, the diagnostic capacity under pressure built on thousands of real incidents, and the trust network with other seniors capable of solving problems that runbooks do not cover. These competences do not fit in a technology list on a CV. They do not pass ATS filters. They have no associated certification.

◆ THE PARADOX OF DECREASING VALUATION

The more tacit knowledge the senior accumulates, the less visible they are in the market. At 45, their certification list is often older than that of a recently trained junior. Their salary is higher — which makes them a cost-reduction target during reorganisations. Their knowledge is deep but invisible — which makes their contribution difficult to justify in a dashboard. The sector thus creates a mechanism of progressive expulsion of its most experienced profiles, precisely when their knowledge reaches its maximum maturity.

◆ NASSIHA — WHAT THIS STUDY DOES NOT CLAIM

This study does not claim that all seniors are victims of discrimination or that all companies deliberately expel their experienced profiles. It identifies structural mechanisms that produce this outcome independently of intentions. An HR director seeking to reduce payroll through budget optimisation who targets the highest salaries is not malevolent — they still produce the expulsion of the senior.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

ALWAYS THE SAME INCIDENT, NEVER TIME TO EXPLORE

The erosion of the senior SRE is not primarily physical. It is cognitive. After 20 years in production, the senior has managed hundreds of incidents on the same systems. They have developed diagnostic patterns of remarkable efficiency — but these patterns are increasingly narrow. They solve very well the problems they have already seen. They solve less well the new problems, because they no longer have time to explore.

Production absorbs everything. On-call covers the nights. Incident tickets fill the days. R&D time is zero or near zero. New technologies arrive — Kubernetes, Rust, eBPF, AI architectures — but the senior cannot practise them. They can read about them, not touch them. They accumulate a growing gap between their theoretical knowledge of new technologies and their operational capability on those same technologies.

◆ THE VICIOUS CYCLE OF HUMAN TECHNICAL DEBT

The more irreplaceable the senior is on legacy systems, the less the organisation can afford to release them to build skills on new systems. The longer they stay on legacy, the more indispensable they become on legacy and the more obsolete on the modern stack. The organisation thus creates a human technical debt — an expert on critical systems who cannot evolve because their evolution would endanger the continuity of the systems they maintain. This debt accumulates silently until the day the senior leaves — and the organisation simultaneously discovers what it had and what it no longer has.

◆ THE SYMMETRY WITH THE BLUE

In "The Blue — What the Diploma Does Not Measure", section 3b documented that the intern brings fresh knowledge that the senior no longer has time to acquire. That observation took the intern's perspective. Seen from the senior's side, the same reality is darker: the senior knows what the intern does not yet know, but they no longer know what the intern already knows. Both need each other. An organisation that separates them — hierarchically, culturally, through recruitment practices — loses both sides of the exchange.

◆ NASSIHA — DESERTIFICATION IS NOT A PERSONAL FAILING

A senior who has not explored new technologies for 5 years is not lazy or resistant to change. They are in a structural situation where exploration time is systematically sacrificed to service continuity. This is not a problem of individual motivation — it is a problem of organisational time allocation. The responsibility belongs to the organisation, not to the individual.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE INCIDENT THAT WOULD HAVE TAKEN 4 HOURS INSTEAD OF 12 MINUTES

When the Last One Who Knows leaves — retirement, resignation, redundancy, burnout — the organisation does not immediately feel what it has lost. In the first weeks, teams manage. The runbooks exist. Monitoring tools alert. Everything seems to work. The loss reveals itself at the first incident no one else has ever seen before.

It is not the routine incident that reveals the senior's absence. It is the rare incident — the one that happens once every three years, on the legacy system that no one has truly understood since they left. This incident takes 4 hours instead of 12 minutes. Or 48 hours instead of 4. Or it does not resolve at all, and an external consultant is called in at €2,000 an hour to urgently reconstruct what the senior knew from memory.

◆ TACIT KNOWLEDGE CANNOT BE DOCUMENTED — IT IS TRANSFERRED

The standard organisational response to a senior's departure is the "handover" — a few weeks of intensive documentation before leaving. This response is structurally insufficient. Tacit knowledge is not the knowledge the senior knows they have. It is the knowledge they mobilise without knowing it — automatic diagnostic patterns, intuitions about abnormal system behaviour, the memory of the contexts in which a decision was made. This knowledge cannot be documented because it is not conscious. It transfers through shared practice, through apprenticeship, through years of working side by side.

◆ THE REAL COST OF DEPARTURE: NEVER CALCULATED

No organisation calculates the true cost of an experienced senior's departure. The cost of replacement is calculated — recruitment, the successor's salary, the training period. What is not calculated is the cost of lengthened incidents, of poorly-made decisions for lack of historical context, of clients lost because a problem was not resolved fast enough, of months spent bringing the successor up to speed on systems they do not yet understand. This invisible cost is systematically higher than the visible cost of recruitment — and nobody sees it coming.

◆ NASSIHA — THIS PROBLEM IS NOT SPECIFIC TO INFRASTRUCTURE

The loss of tacit knowledge when a senior expert leaves is a documented problem across all sectors with a strong technical component — heavy industry, nuclear, aeronautics. Digital infrastructure did not invent this problem. But it worsens it through two characteristics specific to the sector: the speed of technological evolution that accelerates the obsolescence of documented knowledge, and the structurally insufficient documentation culture endemic to the industry.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE EXPULSION MECHANISMS

HOW ORGANISATIONS DISPOSE OF WHAT THEY CANNOT MEASURE

The Last One Who Knows does not always leave voluntarily. Their departure is often the result of a series of organisational mechanisms that progressively push the senior towards the exit — without anyone having explicitly decided to get rid of them.

◆ MECHANISM 1 — THE REORGANISATION THAT DILUTES

Successive reorganisations dissolve the stable teams where the senior built their trust network and informal legitimacy. In a new team with a new manager who does not know their history, the senior starts from scratch in building their credibility — while their years of experience remain invisible to a manager evaluating them on the same criteria as their junior colleagues.

◆ MECHANISM 2 — EVALUATION BY RECENT TOOLS

Annual reviews and skills assessments are built around fashionable technologies. A senior who perfectly masters critical production systems but does not know the latest container orchestration tool will be negatively assessed against a skills framework designed for juniors in training. The evaluation measures what the senior does not yet know — not what they have known for twenty years and that no one else knows.

◆ MECHANISM 3 — THE SALARY AS A TARGET

The senior is the team's most expensive employee. Under payroll pressure, they are the obvious target for budget optimisation. The argument presented is always economic — "reorganisation", "strategic evolution", "rightsizing". The reality is that the senior costs a lot and their cost is visible, while their value is not. The organisation systematically arbitrates in favour of the visible against the invisible.

◆ MECHANISM 4 — BURNOUT BY ACCUMULATION

Twenty years of on-call, twenty years of 3am incidents, twenty years of responsibility for critical systems without that responsibility being recognised at its true value — the senior who does not leave through the economic door often leaves through the exhaustion door. This burnout is not a sudden event. It is a silent accumulation that spreads over years, in a culture that values resilience under adversity without ever questioning the structure that produces that adversity. This mechanism will be examined in depth in Angle 4 — "The Pager and the Body".

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

RETAIN, TRANSMIT, RECOGNISE

There is no miracle solution to the problem of tacit knowledge loss when a senior leaves. There are practices that reduce the risk — and practices that worsen it. Most organisations apply the latter without realising it.

◆ LEVER 1 — STRUCTURAL SENIOR-JUNIOR PAIRING

Tacit knowledge transfer has only one effective vector: shared work on real problems. A senior and a junior working side by side on the same incidents for 12 to 18 months transfer more knowledge than intensive documentation over 3 months before departure. This pairing must be structural — not occasional, not reserved for crisis moments. It must be integrated into the organisation of work as a permanent practice, not an emergency response.

◆ LEVER 2 — GUARANTEED EXPLORATION TIME

The senior's cognitive desertification is avoidable if the organisation guarantees exploration time — protected, regular, non-interruptible by production incidents. 20% of working time dedicated to technological exploration is a documented practice in some organisations. It is uncommon in infrastructure. It should be standard — because a senior who stays current is a senior who stays relevant, and a relevant senior is a senior the organisation has no reason to expel.

◆ LEVER 3 — FORMAL RECOGNITION OF TACIT KNOWLEDGE

The senior must have a formal role that recognises their tacit knowledge and gives them an organisational legitimacy independent of their mastery of recent technologies. Titles like "reference architect", "domain technical expert", or "critical infrastructure lead" are not honours — they are retention and valuation mechanisms that make visible what was invisible. They also protect the senior from reorganisations that reduce experience to a cost and ignore its value.

◆ LEVER 4 — THE DEPARTURE PLAN PREPARED IN ADVANCE

Every experienced senior should have, at any time, a documented transmission plan — not an emergency handover, but a continuous process of documentation and transmission that spans several years. This plan is not a threat ("you are going to leave") — it is an organisational insurance policy ("when you leave, whether in 1 year or 10, the team will not be blind"). This practice is standard in other critical sectors. It is almost absent in digital infrastructure.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

MAKING THE INVISIBLE VISIBLE

The problem of the Last One Who Knows is a visibility problem. Their knowledge is real but invisible. Their cost is visible but their value is not. The solution is not technical — it is organisational and cultural. It consists of making visible what is structurally ignored.

◆ MEASURE 1 — CALCULATE AND PUBLISH THE TRUE COST OF SENIOR DEPARTURES

Every organisation should calculate, after each experienced senior departure, the total cost of that departure — successor recruitment, lengthened incidents in the first 12 months, successor training and support time, external consultants mobilised. This calculation, made visible to management and HR teams, radically changes the cost/value ratio of the senior. A senior at €80,000 annually whose departure costs the organisation €300,000 is no longer expensive — they are underpaid.

◆ MEASURE 2 — INTEGRATE TACIT KNOWLEDGE INTO SKILLS FRAMEWORKS

The skills frameworks used for annual reviews must integrate tacit knowledge criteria: institutional memory, diagnostic capability on legacy systems, transmission to juniors, off-runbook incident resolution. These criteria are no less objective than a certification list — they simply require a different evaluation mode (observation in real situations, peer feedback, incident history). Their absence from current frameworks is a choice, not an inevitability.

◆ MEASURE 3 — CREATE A FORMAL "CRITICAL INFRASTRUCTURE LEAD" STATUS

A formal status recognised in collective agreements, which protects the senior against reorganisations — not for union reasons, but for service continuity reasons. This status is assigned on objective criteria (seniority, scope of responsibility, critical systems mastered), is reviewable, and comes with an obligation of documented transmission. It makes tacit knowledge visible in the organisational chart.

The Last One Who Knows always leaves eventually. The question is not preventing them from leaving. It is ensuring that when they leave, they are not the last one to know.

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART III

LEGAL DOSSIER

9 pieces

HUMAN

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026 · COGNITIVE PSYCHOLOGY

THE ARCHITECTURE OF CONSENT

How Hyperscalers Exploit Cognitive Biases at Each Layer of Perdition

HUMAN

◆ THE CENTRAL THESIS

Hyperscalers did not merely build technical lock-in mechanisms. They — consciously or not — built environments that exploit cognitive biases documented by 70 years of psychology and behavioural economics research. These biases are not individual flaws: they are universal heuristics that become counterproductive in environments specifically designed to instrumentalise them. Each layer of digital perdition has a counterpart in scientific literature. This study maps them — and proposes the counter-heuristics to exit.

PAGES	BIASES	WATERMARK	NUDGES
8	12	HUMAN DOMINANT	4

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

CC BY-NC-SA 4.0

I

SECTION I · THE HOOK — RECIPROCITY AND IDENTITY COMMITMENT THROUGH CERTIFICATION

CIALDINI 1984 · JOULE & BEAUVOIS 1987 — FREE TIER IS NOT GENEROSITY · IT IS PSYCHOLOGICAL DEBT

◆ RECIPROCITY BIAS — Cialdini, Influence, 1984

A received favour creates a psychological obligation to reciprocate — even disproportionately. This mechanism is documented in all human cultures by Cialdini (1984): the beneficiary of a gift feels social and psychological pressure to return, even if the expected return is of incomparable value to the original gift.

AWS Educate offers 750 EC2 hours and $100 credit to the student. The student perceives this gesture as a favour — they feel indebted to the AWS ecosystem for their training and nascent career. When they enter a company, they recommend AWS — not because it is objectively the best choice, but because reciprocity has been activated since university. Free Tier is not marketing — it is an investment in future reciprocity.

◆ IDENTITY COMMITMENT THROUGH CERTIFICATION — Joule & Beauvois, 1987

Commitment theory (Joule & Beauvois, 1987) demonstrates that an individual is psychologically bound to acts they have freely performed. Obtaining an AWS Solutions Architect certification requires months of intensive revision, an exam fee and deep personal investment.

Once certified, the engineer no longer perceives AWS as an interchangeable tool. The certification has modified their professional identity — they are an 'AWS Architect', not an 'infrastructure architect.' Defending a sovereign alternative becomes an existential threat to their hard-won skill capital, validated by the algorithmic ATS sabotage documented in The Newspeak (12p). This identity commitment is deeper than Sunk Cost: it is not financial — it is ontological.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE COGNITIVE SELF-REINFORCEMENT LOOP — THE MORE COMPLEX, THE MORE CREDIBLE

◆ THE IKEA EFFECT — Norton, Mochon & Ariely, Journal of Consumer Psychology, 2012

Individuals systematically overvalue what they have built themselves, regardless of objective quality. A complex cloud architecture built by the internal team over 5 years — Lambda + DynamoDB + API Gateway + SQS + Cognito + CloudFront — is perceived as irreplaceable and sophisticated, not because it is technically superior, but because the team built it. The IKEA Effect creates emotional resistance to migration that has nothing to do with the real TCO.

◆ THE DUNNING-KRUGER EFFECT OF THE IAC SCRIPT — Kruger & Dunning, JPSP, 1999

The Dunning-Kruger effect demonstrates that the least competent individuals in a domain massively overestimate their mastery. The DevOps merge and IaC give the developer the perfect illusion that writing 20 lines of HCL equates to mastering infrastructure. They do not see ARP tables, network segments, inodes, MCO — because they believe they master them via the script. They do not see the lock-in trap because they believe they are taming the machine. Low-level incompetence is invisible to those who suffer from it — that is precisely the definition of Dunning-Kruger.

◆ THE SELF-REINFORCEMENT LOOP — PROPORTIONALITY BIAS (Rao & Monroe, 1989)

The tendency to perceive quality proportionally to complexity and cost. A cloud architecture stacking 6 managed services is perceived as 'enterprise-grade' and professional — precisely because it is complex. A simple, reliable bare-metal server is perceived as 'artisanal.' Hyperscaler artificial complexity is a perceived quality signal — the more unreadable the architecture, the more serious it seems. The loop: complexity → dopaminergic satisfaction → IKEA overvaluation → dismantling resistance → more complexity.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE SEMANTIC TRAP — AVAILABILITY · HALO · SOCIAL CONTAGION

KAHNEMAN 1973 · THORNDIKE 1920 · CHRISTAKIS & FOWLER 2009 — NEWSPEAK AS ARCHITECTURE OF PERCEPTION

◆ AVAILABILITY HEURISTIC — Tversky & Kahneman, Cognitive Psychology, 1973

Decisions are influenced by the ease with which examples come to mind. 'Serverless', 'Cloud-Native', 'AI', 'Digital Transformation' are omnipresent in conferences, LinkedIn, tech media and Gartner reports. Their cognitive availability makes them perceived as normative — what everyone does, what is modern. The bare-metal alternative has no AWS re:Invent conference. No LinkedIn coverage. Cognitively unavailable — therefore perceived as marginal, even when technically superior.

◆ THE HALO EFFECT — Thorndike, Journal of Applied Psychology, 1920

A positive characteristic projects a positive aura onto unrelated domains. Google = brilliant search engine → Google Cloud = necessarily excellent. Amazon = global logistics → AWS = reference infrastructure. Brand prestige replaces technical evaluation. The board member who orders on Amazon Prime perceives AWS differently from one who has only used AWS infrastructure services. The halo effect is the cognitive shortcut that avoids objective evaluation.

◆ SOCIAL CONTAGION OF HYPE — Christakis & Fowler, Connected, 2009

Christakis & Fowler documented that behaviours, opinions and emotions spread through social networks up to 3 degrees of separation. When all Paris CTOs adopt AWS and discuss it on LinkedIn, this adoption spreads through their networks as a behavioural norm. The resister is perceived as an outlier — not as a rigorous engineer. Social contagion pressure replaces individual technical judgement. This mechanism explains the adoption speed of the 'Mystical Machine' (cloud AI) in 2023-2024: not a rational decision — an accelerated social contagion.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

FESTINGER 1957 · JANIS 1972 · THALER & SUNSTEIN 2008 — DEFAULT CHOICE AS RESPONSIBILITY DILUTION

◆ SUNK COST + COGNITIVE DISSONANCE — KAHNEMAN & TVERSKY 1979 · FESTINGER 1957

The Sunk Cost Fallacy (documented in The Free Sample S4) drives continued commitment because one has already invested — even when irrational. Cognitive dissonance (Festinger, 1957) adds a layer: when new information conflicts with past behaviour, the individual prefers to reject the information rather than acknowledge the error. The CIO who evangelised cloud for 5 years cannot integrate corpus data — they create unbearable dissonance. They do not reject them because they are incompetent: they reject them because accepting them would threaten their psychological coherence.

◆ DEFAULT EFFECT AS RESPONSIBILITY DILUTION — THALER & SUNSTEIN 2008 · DARLEY & LATANÉ 1968

The default option is chosen because it avoids decision effort (Thaler & Sunstein, 2008). Modern frameworks (Next.js, Spring Boot) and documentation propose managed cloud deployment as the default configuration. Lock-in installs without deliberate decision — through differential friction between the proprietary option (zero effort) and the neutral option (engineering effort).

But the deepest dimension is responsibility diffusion (Darley & Latané, 1968): choosing the AWS + Gartner + McKinsey default option dilutes individual responsibility to zero. The decision-maker unconsciously thinks: 'If I choose the option recommended by everyone, I cannot be blamed for an outage.' Choosing custom Bare-Metal = total individual responsibility. An unconscious career protection strategy as much as a cognitive mechanism.

◆ GROUPTHINK — JANIS, VICTIMS OF GROUPTHINK, 1972

In a cohesive group under pressure, critical thinking is suppressed in favour of consensus. When all CTOs go on AWS, the one proposing Bare-Metal is perceived as aberrant — not as a rigorous engineer. Conformity pressure replaces technical judgement. Groupthink requires no bad intention — it operates mechanically as soon as consensus is perceived as the norm.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

KAHNEMAN & TVERSKY 1979 · OMER & ALON 1994 · SAMUELSON 1988 — BIASES THAT MAKE INACTION PREFERABLE TO ACTION

◆ LOSS AVERSION — Kahneman & Tversky, Econometrica, 1979

The pain of a loss is psychologically twice as intense as the pleasure of an equivalent gain. 'Not running out of cloud capacity' feels more urgent than 'paying half as much.' Organisations massively over-provision on cloud — and never calculate what well-sized bare-metal would cost — because the risk of shortage is felt more strongly than the cost of excess. Hyperscalers exploit this bias with Pay-as-you-go pricing: 'You only pay for what you consume' hides 'You never know how much you will pay.'

◆ NORMALCY BIAS — Omer & Alon, 1994

The tendency to underestimate the probability and impact of rare but catastrophic events. A Taiwan blockade, a CLOUD Act activation on sensitive data, a 24-hour regional AWS outage — all individually improbable, all with documented precedents (The Taiwan Bottleneck · The Hand Over Nations). Normalcy bias explains why organisations do not build continuity plans without cloud: 'it won't happen to us.' The psychological complement to the Taiwan Bottleneck — the cognitive mechanism that makes geopolitical risk theoretical even when documented.

◆ STATUS QUO BIAS — Samuelson & Zeckhauser, Journal of Risk and Uncertainty, 1988

The preference for the current state of things, even when change would be beneficial. Once infra teams have been dissolved (NoOps), nobody proposes recreating them — not because it is irrational, but because the status quo has its own psychological inertia. The anticipated pain of change (rebuilding skills, reconstructing teams, justifying cost) is overestimated relative to future benefits — exactly like loss aversion, but applied to the entire organisation.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · LIBERATION NUDGES — FOUR OPERATIONAL COUNTER-HEURISTICS

TURNING BIASES AGAINST THEMSELVES — FROM AUDIT TO ACTION

◆ NUDGE 1 · INVERTING LOSS AVERSION — Tversky & Kahneman, Science, 1981 (Reframing)

Do not promise an abstract future gain (sovereignty) — show an immediate, quantifiable loss. Cognitive reframing (Tversky & Kahneman 1981) demonstrates that the same decision presented as 'avoided loss' rather than 'obtained gain' produces radically different decision responses.

Application: Calculate cumulative egress fees over 10 years from existing billing data. Show the current exit ransom cost. Present Bare-Metal not as 'going backwards' but as 'elimination of a documented recurring loss.' Dropbox saved $75 million over 2 years repatriating servers (2017, publicly documented). Basecamp saves over $1.5 million per year since its cloud exit (2023, David Heinemeier Hansson). Not exceptions — industrial margin optimisation decisions.

◆ NUDGE 2 · THE EXIT NARRATIVE WITHOUT LOSS OF FACE

The 2019 decision was correct with 2019 information. 2026 information is different — CLOUD Act clarified, egress fees documented, TSMC risk identified, Schrems II published. Revising one's position in light of new data is not capitulation — it is the definition of good management. This narrative deactivates cognitive dissonance by transforming the position change into an act of managerial competence rather than an admission of error.

◆ NUDGES 3 & 4 · INVERTED CONFORMITY + ALTERNATIVE AVAILABILITY

Nudge 3 — Inverted conformity: show that Dropbox, Basecamp, and dozens of profitable organisations have left the cloud. The return to Bare-Metal becomes a mark of distinction and maturity, not a delay. Reverse social contagion in the direction of sovereignty — if tech profitability leaders do it, it is the new norm.

Nudge 4 — Make the alternative cognitively available: the Operation Dindon corpus is the answer to the availability heuristic. Sixty-four structural studies = sixty-four times the sovereign alternative is made cognitively available. Each published study increases the cognitive availability of the bifurcation.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

7

CLOSING · SYSTEM 1, SYSTEM 2 — THE UNIFYING THESIS

KAHNEMAN, THINKING FAST AND SLOW, 2011 — EACH LAYER WAS DESIGNED FOR SYSTEM 1

Kahneman distinguishes two thinking systems: System 1 (fast, intuitive, heuristic — operates without effort) and System 2 (slow, deliberate, analytical — requires conscious cognitive effort). Hyperscalers designed their acquisition journey to maximise System 1 at each friction point: Free Tier activates without reflection, the default option is accepted without deliberation, newspeak is received as evidence, AI hype spreads through social contagion.

System 2 — reading the Terms, calculating real TCO, the architectural question 'outside AWS without rewrite?', identifying responsibility diffusion in the Default Effect — requires deliberate effort that the acquisition journey was designed never to trigger. This is the corpus's unifying thesis: the 64 Operation Dindon studies do the work of System 2 on behalf of the decision-maker. They make visible what biases render invisible.

These biases are not flaws — they were instrumentalised.

Not individual weaknesses — universal heuristics

placed in environments designed to make them counterproductive.

Resistance is not suspicion — it is System 2.

Amine RAITI · Operation Dindon · 2026

NEMO SUPRA LEGEM EST

HUMAN

CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

TERMS OF SERVICE UNDER THE MICROSCOPE

AWS & GCP — Corpus Reading

What the Terms of Service Say That the Brochures Do Not

RATIO

◆ METHODOLOGY

This study analyses the Terms of Service of AWS (June 2026 version) and Google Cloud Platform in light of the Opération Dindon corpus. Each identified clause is paraphrased — never quoted verbatim — and placed alongside the corpus study that documents its effect. The objective is not legal. It is to show that what the corpus named — cognitive capture, infrational loop, lock-in, loss of sovereignty — is literally inscribed in the contractual documents that European organisations sign every day.

DOCUMENTS	CLAUSES ANALYSED	WATERMARK
AWS + GCP	**8**	**RATIO**

Amine Raiti — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

EUROPEAN DATA UNDER US LAW — THE TERMS OF SERVICE SAY SO EXPLICITLY

◆ WHAT GCP SAYS — CALIFORNIA JURISDICTION (SECTION 14.12)

The Google Cloud Platform Terms of Service stipulate, for all entities that are not US government entities, that the law of the State of California governs the contract — excluding California's conflict of laws rules — and that the competent courts are those of Santa Clara County, California. The European customer who signs the GCP Terms therefore accepts that any dispute be resolved under California law, by Californian courts.

What this means concretely: the data of a French hospital, a Belgian bank or a German public authority hosted on GCP are contractually subject to US law. The GDPR applies to collection and processing — but the contractual reference framework is the law of the State of California.

◆ WHAT AWS SAYS — ACCESS UNDER GOVERNMENT ORDER

The AWS Terms of Service specify that the company may act in accordance with any applicable judicial, regulatory or governmental order or request. In the event of content removal without prior notice, AWS will notify the customer promptly — unless prohibited by law. This formulation — "unless prohibited by law" — is precisely the window of the CLOUD Act and National Security Letters: a non-disclosure order (gag order) is legally possible, and the Terms of Service anticipate it.

◆ CORPUS LINK — "THE DIGITAL GROUND" · "NO SOVEREIGNTY WITHOUT MATTER"

The Opération Dindon corpus documented in "The Digital Ground" that AWS and GCP European datacentres do not create European sovereignty — because the applicable jurisdiction for data remains American. This study demonstrates that this is not a corpus interpretation — it is textually inscribed in the Terms of Service. The California jurisdiction clause of GCP and the government access clause of AWS are the contractual proof of the Digital Ground thesis.

The 8 September 2026 ultimatum demands resiliable commits, reasonable egress fees, real portability, and the CLOUD Act limited to its territory. The California jurisdiction clause demonstrates that the fourth demand is the most fundamental: without a change of jurisdiction, the other three are insufficient.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THEY CAN CHANGE PRICES AND CONDITIONS — AT ANY TIME

◆ WHAT GCP SAYS — FEE REVISIONS (SECTION 2.6)

Google reserves the right to change fees at any time, unless otherwise expressly agreed in an addendum or Order Form. For GWS, Looker and Cloud Identity services, Google commits to 30 days advance notice. But for GCP — the core cloud services — no notice period is mentioned. The formulation "at any time" is explicit.

This unilateral modification of pricing is also conditional: if the customer disagrees with a GCP pricing change, they may terminate for convenience — but under the termination conditions described elsewhere, which include deadlines and obligations.

◆ WHAT AWS SAYS — RESERVED INSTANCE MODIFICATIONS

AWS states it may change Savings Plans, EC2 Reserved Instances and Dedicated Host Reservation pricing at any time — but price changes do not apply to already-purchased reservations. However, if Microsoft increases Windows licence fees or Red Hat increases RHEL fees, AWS may pass these increases on with 30 days notice, including to existing Reserved Instances.

The cascade mechanism: AWS is not the only variable. Its own suppliers' (Microsoft, Red Hat, Oracle) price increases can be passed to the end customer without that customer having any recourse up the chain.

◆ WHAT AWS SAYS — NON-CANCELLABLE COMMITS

For Savings Plans, EC2 Reserved Instances, Capacity Blocks, DB Savings Plans and SageMaker AI Savings Plans, AWS Terms specify that commitments are non-cancellable and that the customer will be charged for the full duration of the selected term, even if the agreement is terminated. Amounts paid are non-refundable, unless AWS terminates the agreement other than for cause.

◆ CORPUS LINK — FIRST ULTIMATUM DEMAND: "RESILIABLE COMMITS"

The first of the four 8 September 2026 demands is: resiliable commits. AWS Terms document precisely the opposite: commitments are "noncancellable". The corpus thesis is verified by the contractual text. This is not an opinion — it is a clause.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THEY CAN CUT — YOU CANNOT LEAVE

◆ UNILATERAL SUSPENSION — AWS AND GCP

AWS: AWS may suspend or terminate access to any Beta service at any time, without notice. For generally available services, suspension is possible on governmental or judicial order, or if AWS judges that the security of its services is threatened. Depending on the service, termination may occur with varying notice periods.

GCP: Google may suspend all or part of access to the services if it reasonably believes suspension is needed to protect the services, other customers or third parties. Google may also suspend immediately to comply with any applicable Legal Process. The suspension may precede notification — if the law permits.

◆ DATA AFTER TERMINATION — THE 30-DAY WINDOW

AWS (Supply Chain N-Tier): when the service ends, the customer has 30 days following termination to retrieve their content. After this deadline, the content will be removed. This formulation — present in several AWS Terms sections — is the migration window: 30 days to export years of structured data, in a format that may not be interoperable with the target system.

GCP: upon termination, all rights to access the services (including access to Customer Data, if applicable) cease. The precise conditions for post-termination data retention and export refer to the Service Specific Terms — a separate document not integrated into the general Terms analysed here.

◆ GCP — TERMINATION FOR CONVENIENCE BY GOOGLE (Section 8.5)

For GCP and TSS services only, Google may terminate the agreement or any applicable Order Form for its own convenience at any time with 30 days prior written notice. The customer may also terminate for convenience at any time on written notice — but must immediately cease using the services upon termination. The asymmetry is structural: Google can terminate for convenience with 30 days notice. The customer can too — but immediately loses access to data and services.

RAITO

◆ CORPUS LINK — "THE GENTLE EXIT" · "THE SOVEREIGN INTERFACE"

The Gentle Exit proposed in the corpus describes progressive migration, service by service. The 30-day Terms window says the opposite: exit is necessarily abrupt if imposed by the provider. The Sovereign Interface documented in the corpus is precisely the architectural response to this constraint: decouple application code so that migration is not an emergency.

Amine RAITI · CC BY-NC-SA 4.0

YOU ASSUME — THEY LIMIT — YOU DEFEND — THEY ARE DEFENDED

◆ LIABILITY LIMITATION — AWS AND GCP

AWS: the aggregate liability of AWS and its affiliates for Beta services is limited to the amount actually paid by the customer during the 12 months preceding the claim. For generally available services, the same logic applies with similar caps per Customer Agreements. Indirect, consequential, special or punitive damages are excluded.

GCP: neither party will have liability for indirect, consequential, special, incidental or punitive damages, or for lost revenues, profits, savings or goodwill. The limitation applies to the extent permitted by applicable law. Aggregate amounts are capped per the Service Specific Terms.

The asymmetry: if an AWS or GCP failure causes €10 million in lost business for a customer organisation, the provider's contractual liability will be limited to fees paid over 12 months — perhaps €50,000. The gap between actual loss and contractual compensation is structural.

◆ ASYMMETRIC INDEMNIFICATION — THE CUSTOMER DEFENDS THE PROVIDER

GCP: the customer agrees to defend Google and its affiliates against any third-party legal proceeding arising from the customer's non-compliant use of the Terms, or from a claim that customer data infringes third-party intellectual property rights.

AWS: Marketplace sellers indemnify AWS and its affiliates against any transaction-related tax claims. More generally, the customer is responsible for all use under their account, including by their end users.

The structure: Google defends the customer against claims concerning Google's services (Google's intellectual property). The customer defends Google against claims concerning customer usage. Usage liability migrates to the customer. Platform liability stays with Google but is capped.

◆ CORPUS LINK — "THE STATE AND THE INVISIBLE BODY" · "NO SOVEREIGNTY"

The corpus documents the absence of infrastructure body protection by the State. This clause shows that protection does not exist in the contract either: provider liability is capped, customer liability on usage is unlimited. The sovereign organisation is the one that does not depend on an asymmetric contract for its business continuity.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

CERTIFICATIONS AS A DEPENDENCY TOOL · PORTABILITY ABSENT

◆ THE AWS CERTIFICATION PROGRAMME — THE INFRATIONAL LOOP IN THE CONTRACT

The AWS Terms include a dedicated article on the "AWS Certification Program" — a service in its own right subject to a separate agreement (Certification Program Agreement, CPA). AWS certification certifies that a person has passed a proficiency test on a specific AWS service or domain. It is not transferable between providers and does not guarantee the quality of services delivered by the certified provider.

What this says about the job market: an AWS certification is an employable asset on the AWS market — and only on that market. An AWS Certified Solutions Architect has an asset worth something on AWS projects. There is no equivalent certified on GCP, Azure, or on bare-metal. A CPF budget funding an AWS certification funds dependency — not neutral competence.

What the Terms confirm: the certification programme is a "Service" under the Terms of Service. AWS defines its conditions, may modify programme terms, and the CPA supersedes the general agreement in case of conflict. The certification belongs to the AWS ecosystem — not to the engineer.

◆ DATA PORTABILITY — THE REVEALING ABSENCE

Neither the AWS Terms analysed, nor the GCP Terms analysed, contain a clause guaranteeing data portability in a standard, interoperable format, at no additional cost. Searching for "portability" in both documents returns only:

— AWS: telephone number portability in Amazon Chime and Amazon Connect (telephony portability, not cloud data portability).

— GCP: HIPAA (Health Insurance Portability and Accountability Act) — a US healthcare law, not a cloud data portability guarantee.

The absence is the data point: in 302,000 characters of AWS Terms and 113,000 characters of GCP Terms, the term "data portability" in the sense of a guaranteed export in a neutral format at no additional cost is absent. These are the second and third ultimatum demands — "real portability" and "reasonable egress fees" — that the Terms do not address.

◆ CORPUS LINK — "THE CPF MISUSED" · "THE INFRATIONAL LOOP" · ULTIMATUM

The CPF Misused documents that training budgets flow to proprietary cloud certifications. The AWS Terms confirm that these certifications are AWS services, governed by AWS law, terminable by AWS. The loop is complete: public money funds competences that belong to the ecosystem of an extraterritorial actor.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

CORRESPONDENCE TABLE — TERMS CLAUSE × CORPUS STUDY × ULTIMATUM DEMAND

TERMS CLAUSE	WHAT IT SAYS	CORPUS STUDY	ULTIMATUM DEMAND
GCP §14.12 — California law	Any dispute resolved under California law, Santa Clara courts	The Digital Ground · No Sovereignty Without Matter	**4. CLOUD Act limited to its territory**
AWS — Government access	Disclosure possible under government order, no notice if law prohibits	The Digital Ground · Anatomy of the Loss	**4. CLOUD Act limited to its territory**
GCP §2.6 — Fees "at any time"	Google may change GCP fees at any time with no guaranteed notice	The Infrational Loop · FinOps of the Last Gram	**2. Reasonable egress fees**
AWS — "Noncancellable" commits	Reserved Instances, Savings Plans: billed even if agreement is terminated	Anatomy of the Loss · The Gentle Exit	**1. Resiliable commits**
AWS/GCP — Unilateral suspension	Suspension possible at any time, immediately for "Legal Process"	The Digital Ground · The State and the Invisible Body	**1 + 4. Resiliable commits + CLOUD Act**
AWS/GCP — Data: 30 days post-termination	30 days to retrieve data — then deletion. No neutral format guarantee.	The Gentle Exit · The Sovereign Interface	**3. Real portability**
AWS/GCP — Liability limitation	Liability capped at last 12 months' fees. Indirect damages excluded.	The Economy of Silence · The State and the Invisible Body	**All 4 demands**
AWS — Certifications as proprietary service	AWS Certification Programme is an AWS service under a separate AWS agreement	The CPF Misused · The Infrational Loop	**3. Real portability (of competences)**
AWS/GCP — Portability absent	No clause guaranteeing data export in a neutral format at no additional cost	No Sovereignty · The Sovereign Interface	**2 + 3. Egress fees + Portability**

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THIS IS NOT AN OPINION — IT IS IN THE CONTRACT YOU SIGNED

◆ WHAT THIS ANALYSIS DEMONSTRATES

The Opération Dindon corpus named phenomena — cognitive capture, infrational loop, loss of sovereignty, lock-in, extraterritorial dependency. These phenomena could have been contested as interpretations, opinions, or rhetorical exaggerations. This analysis of the Terms of Service demonstrates that they are textually inscribed in the contractual documents that European organisations sign every day:

— GCP's California jurisdiction is in Section 14.12.
— The unilateral fee modification "at any time" is in GCP Section 2.6.
— "Noncancellable" commits are in the AWS Reserved Instances, Savings Plans and Capacity Blocks sections.
— Suspension on "Legal Process" without mandatory prior notice is in both Terms.
— Data portability in a neutral format is absent from both documents.
— Liability limitation to last 12 months' fees is in both Terms.
— The AWS Certification Programme is a proprietary service governed by a separate agreement.

◆ THE FOUR 8 SEPTEMBER 2026 DEMANDS — JUSTIFIED CLAUSE BY CLAUSE

1. Resiliable commits: justified by the "noncancellable" clauses in AWS Savings Plans, Reserved Instances, Capacity Blocks, DB Savings Plans.

2. Reasonable egress fees: justified by the absence of any clause governing outbound transfer fees, and by the possible fee revision "at any time" in GCP.

3. Real portability: justified by the absence of any clause guaranteeing data export in a neutral format at no additional cost, and by the 30-day post-termination window.

4. CLOUD Act limited to its territory: justified by GCP's California jurisdiction clause (Section 14.12) and AWS's government access clause that anticipates legal non-disclosure orders.

The Opération Dindon corpus did not invent these problems.
It named them before you read the Terms of Service.
The Terms of Service confirm what the corpus said.
The hyperscalers' silence does too.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

WORKING PAPER · SYSTEMIC ANALYSIS · JUNE 2026

LOCK-IN BY DESIGN

A Six-Layer Taxonomy of Retention Mechanisms in Hyperscaler Cloud Contracts

STRUCTURED ABSTRACT

BACKGROUND

Amazon Web Services, Microsoft Azure and Google Cloud Platform collectively hold over 65% of the global cloud infrastructure market (Synergy Research Group, Q1 2026). The academic literature on cloud vendor lock-in has predominantly addressed technical dimensions — proprietary APIs and non-portable data formats. Contractual, linguistic and behavioural dimensions have received comparatively little formal treatment.

RESEARCH GAP

Three of the six lock-in layers identified in this study — Skills Lock-in (Layer 4), Demand Lock-in (Layer 5) and Cognitive Lock-in (Layer 6) — are, to the author's knowledge, absent from the formalised academic literature on cloud computing. The mechanism of linguistic substitution as a vector for legal rights waiver is likewise undocumented in existing works.

METHODS

Comparative analysis of 47 versions of contractual documents (General Terms, Product Terms, DPAs, SLAs and Addenda) published by AWS, Google Cloud and Microsoft Azure in French, English and German between 2022 and 2026. Algorithmic diff combined with manual legal review against French, German, UK and EU law.

FINDINGS

Six distinct and cumulative lock-in layers are identified. A specific undocumented mechanism is evidenced: the linguistic substitution clause, whereby the client-language version is declared "informational only" while the English version "prevails", enabling the inclusion of rights-waiver clauses in the English-only version absent from the client-language version. The temporal dimension is modelled in three phases (Adoption · Dependency · Captivity) with a formalised economic point-of-no-return.

POLICY IMPLICATIONS

Four regulatory proposals are formulated, including the extension of EU Data Act Article 25 to pricing-based switching obstacles (egress fees) and the formal designation of competent Data Act authorities by EU Member States.

Keywords: vendor lock-in · cloud computing · unfair contract terms · digital sovereignty · EU Data Act · data portability · linguistic substitution · switching costs · platform regulation

AMINE RAITI

Infrastructure Architect & SRE · Paris · Independent researcher

This document is based on the author's independent analysis of publicly available contractual documents.
It does not constitute legal advice. The author has no affiliation with the companies mentioned.

Public document · CC BY-NC-SA 4.0

Opération Dindon

I. INTRODUCTION & METHODOLOGY

1.1 Context and motivation

Cloud infrastructure market concentration has reached historically high levels. According to data published by Synergy Research Group for Q1 2026, Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP) jointly account for 67% of the global cloud services market. This concentration creates structural dependency for private and public organisations on an oligopoly whose contractual practices have attracted increasing regulatory scrutiny — including the Bundeskartellamt's §19a GWB proceedings against Amazon (2022) and Microsoft (2023), the Competition and Markets Authority's Cloud Market Study (2023), and the entry into force of the EU Data Act (Regulation (EU) 2023/2854, applicable since September 2025).

The academic literature on cloud vendor lock-in (Armbrust et al., 2010; Petcu, 2013; Opara-Martins et al., 2016) has predominantly addressed technical dimensions of lock-in — proprietary APIs, non-portable data formats, single-vendor architectures. Economic dimensions (Farrell & Klemperer, 2007) have been partially addressed in the form of switching cost analysis. Contractual, linguistic and behavioural dimensions remain underdocumented. This study aims to address these gaps through a systematic analysis of publicly available contractual documents.

1.2 Methodology

The analysis rests on the comparative examination of **47 versions of contractual documents** published by AWS, Google Cloud and Microsoft Azure over the period 2022–2026, including General Terms of Service, Product Terms, Data Processing Agreements (DPAs), Service Level Agreements (SLAs), and specific addenda (Data Act Addendum, UK Switching Addendum, EES, MOSA). Documents were collected in French, English and German where available.

◆ CORPUS

47 contractual documents · 3 hyperscalers · 3 languages (FR/EN/DE) · Period 2022–2026 · Including: 19 ToS versions · 14 Product Terms · 8 DPAs/SLAs · 6 specific addenda. Full corpus in Annex A.

◆ METHOD

Normalised textual diff (section by section) + manual legal review of substantive modifications. Confrontation against French law (Civil Code, Loi Toubon), German law (BGB, GWB), UK law (UCTA 1977, CRA 2015, CA 1998) and EU law (Data Act, GDPR). Methodology detailed in Annex D.

1.3 Scope and limitations

The analysis covers exclusively publicly available standard contractual documents. Individually negotiated contracts (Enterprise Agreements, sector-specific agreements) may differ significantly. Microsoft Azure's monthly modification of Product Terms — documented at 608 occurrences between January and June 2026 — creates documentary instability that limits the temporal scope of conclusions. Results should be read as an analysis of the standard public conditions applicable to the majority of non-negotiating clients, not as an exhaustive description of all possible contractual configurations.

1.4 Structure

Section 2 presents the six-layer taxonomy. Section 3 analyses the linguistic substitution mechanism. Section 4 models the temporal dimension of lock-in. Section 5 examines competition law implications. Section 6 sets out research and regulatory proposals. Annexes A–D reproduce the primary data underlying the analysis in full.

[1] *Synergy Research Group, "Cloud Market Share Q1 2026", published April 2026.*
[2] *Armbrust M. et al., "A View of Cloud Computing", Communications of the ACM, 53(4), 2010, pp. 50–58.*
[3] *Farrell J. & Klemperer P., "Coordination and Lock-In: Competition with Switching Costs and Network Effects", Handbook of Industrial Organization, Vol. 3, Elsevier, 2007.*
[4] *Opara-Martins J. et al., "Critical analysis of vendor lock-in and its impact on cloud computing migration", JCSA, 5(4), 2016.*

Amine RAITI · CC BY-NC-SA 4.0

2. TAXONOMY — FROM VISIBLE TO ENTRENCHED

2.1 Positioning within existing literature

The classical definition of vendor lock-in (Shapiro & Varian, 1999[5]) is grounded in switching costs — the costs incurred when transitioning from one supplier to another. Porter (1980)[6] distinguishes barriers-to-exit (retaining the customer in the relationship) from barriers-to-entry (preventing a competitor from entering). Farrell & Klemperer (2007) provide the most comprehensive treatment of switching costs in platform economics, identifying installed-base effects, learning costs and contractual switching costs as the primary categories.

The taxonomy proposed here extends these frameworks by identifying three additional layers — behavioural (Skills, Layer 4), decisional (Demand, Layer 5) and cognitive (Layer 6) — which are absent from existing cloud lock-in literature and which, empirically, may be the most durable obstacles to supplier switching in mature cloud deployments.

2.2 The six layers — formal definitions

◆ LAYER 1 · CONTRACTUAL LOCK-IN — SURFACE · VISIBLE WITH CAREFUL READING

Definition: contractual provisions that create an asymmetry of rights between supplier and client, making termination costly, complex or legally uncertain. **Documented mechanisms:** (a) linguistic substitution clause — the client-language version is declared "informational only" while the English version "prevails"; (b) selective non-cancellability — certain commitments survive termination of the master agreement (AWS s.5.4.2; Azure EES Savings Plans); (c) automatic unilateral modification — Product Terms may be modified without individual notice (Azure: 608 modifications Jan.–Jun. 2026); (d) cascade incorporation by reference — the most constraining clause is systematically located several layers below the signed document.

◆ LAYER 2 · PRICING LOCK-IN — SUB-SURFACE · ABSENT FROM CONTRACT AT SIGNING

Definition: pricing structure creating an asymmetry between entry cost (zero or negative: onboarding credits) and exit cost (positive and growing with data volume). **Documented mechanisms:** egress fees ($0.07–0.12/GB to the Internet depending on provider and region); early termination charges (Azure: 12% of residual commitment + $50,000/yr cap); absence of contractual caps on egress fees in public Terms of Service.

◆ LAYERS 3–6 · SUMMARY — MID-DEPTH TO BEDROCK

Layer 3 — Technical: proprietary APIs without standardised equivalent (e.g. BigQuery SQL dialect) · managed services without native portability · superlinear effect: migration cost of n interconnected services > n × cost of migrating one. **Layer 4 — Skills (original concept):** subsidised proprietary certification programmes creating professional attachment bias; passive resistance from middle management as an unconscious sabotage vector in migration projects. **Layer 5 — Demand (original concept):** business units formulate their needs as named products ("Azure OpenAI") rather than as functional requirements, closing competitive evaluation before technical assessment begins. **Layer 6 — Cognitive (original concept):** after 3–5 years of use, organisational KPIs are defined within the supplier's console; any migration implies the loss of the historical performance benchmark.

[5] *Shapiro C. & Varian H.R., Information Rules: A Strategic Guide to the Network Economy, Harvard Business School Press, 1999.*
[6] *Porter M., Competitive Strategy: Techniques for Analyzing Industries and Competitors, Free Press, 1980.*
[7] *Petcu D., "Portability and Interoperability between Cloud Providers: Challenges and Case Study", Towards a Service-Based Internet, LNCS 8135, Springer, 2013.*

 Amine RAITI · CC BY-NC-SA 4.0

3. LINGUISTIC SUBSTITUTION AS A VECTOR FOR RIGHTS WAIVER

3.1 Description of the mechanism

All three hyperscalers publish their Terms of Service in multiple languages while stipulating that the English version "prevails" or "governs" in case of divergence. This mechanism, ostensibly a practical translation disclaimer, produces three distinct legal effects: (a) it deprives the client of certainty that the version in their own language is binding on the supplier; (b) it allows the supplier to publish updates in English without awaiting translation into other languages, creating temporal gaps during which the client operates under an outdated version; (c) in the most significant cases documented, it allows clauses waiving statutory legal rights to be included in the English-only version, absent from the client-language version.

This mechanism is rendered legally significant by mandatory language legislation applicable in several EU jurisdictions. In France, the *Loi Toubon* (Law n°94-665 of 4 August 1994)[8] requires that contracts concluded with French legal entities for activities conducted in France be drafted in French, the French version being legally binding. A contractual clause declaring the English version to prevail is contrary to this mandatory provision and is therefore void *ab initio* under French law, without any requirement for the client to have challenged it at signing.

3.2 Measured translation delays — primary data

Provider	Linguistic strategy	FR delay	DE delay	Missing updates
AWS	Explicit clause "EN prevails" in header of each page	**37 days**	**37 days**	Variable per version
GCP	Concealed clause s.14.18 · FR version frozen Aug. 2023	**2 yrs 9 mths**	**4 years**	**FR: 13 · DE: 16**
Azure EES	Self-admitted: "Published Feb 1 EN, translated Feb 9"	**8 days (admitted)**	**8 days (admitted)**	Capacity Blocks Jun. 2026: 0 days

3.3 Documented case — waiver of Directive 2018/1972 rights (Azure EES)

◆ RIGHTS WAIVER CLAUSE — ABSENT FROM FRENCH VERSION | MOST SIGNIFICANT CASE DOCUMENTED

The Microsoft EES Product Terms (January and June 2026, English versions) contain the following clause, absent from the French version:

> *"Customer agrees to waive any and all entitlements that would otherwise be applicable under the European Electronic Communications Code (Directive 2018/1972) Article 102 paragraphs 1, 3, and 5 ; Article 105 paragraph 1 ; and Article 107 paragraphs 1 and 3."*

The EES is the standard contractual vehicle through which public institutions — including educational bodies — subscribe to Microsoft services under public procurement frameworks. Any organisation operating under EES terms is potentially bound by this waiver. The waiver of statutory rights derived from a public-order European directive — effected contractually and presented in a single language — raises, under the analysis conducted, issues of compliance with: the French *Loi Toubon* art. 5; the French Civil Code art. 1171 (significant imbalance in standard-form contracts); and Directive 2018/1972 itself, whose provisions are mandatory in nature and cannot be validly waived by private contract.

3.4 Regulatory implications

Beyond French mandatory language law, the EU Data Act art. 23 requires that contractual terms be provided in a manner that is "easy to read, understand and accessible". The structural coexistence of an "informational" client-language version and a legally binding English version directly contradicts this accessibility requirement. In German law, BGB §305c provides that terms which, in the circumstances, a contracting party could not reasonably expect, are not incorporated into the contract — a provision directly applicable to clauses present only in the English version of a contract signed by a German entity.

[8] *Loi n°94-665 du 4 août 1994 relative à l'emploi de la langue française (Loi Toubon), art. 5. English translation: "Any document containing obligations for a natural or legal person in the context of the performance of a public service mission must be drafted in French."*

[9] *Directive (EU) 2018/1972 of 11 December 2018, establishing the European Electronic Communications Code, OJ L 321, 17 December 2018.*

[10] *Regulation (EU) 2023/2854 (Data Act), art. 23: "Data processing services contracts shall be drawn up in a clear and comprehensible manner."*

Amine RAITI · CC BY-NC-SA 4.0

4. PROGRESSIVE DEEPENING OF LOCK-IN OVER TIME

4.1 Accumulation model — three phases

Cloud lock-in does not crystallise at the moment of signing but deepens progressively through use. We propose a three-phase model, calibrated on medium-sized organisations (100–1,000 employees, cloud spend $200k–$2M/year):

◆ PHASE 1 · ADOPTION (YEAR 1) — Estimated lock-in score: 2/10 — ACTION WINDOW OPEN

Contractual commitments are typically short (pay-as-you-go or 1-year Reserved Instances). Data volumes are limited, rendering egress fees immaterial. The number of proprietary services adopted is low (2–3), and migration remains technically feasible in a matter of weeks. Teams are in the process of obtaining certifications — the attachment bias is not yet established. Business units have not yet entrenched the pattern of named-product requests. **Recommendation: this is the only phase in which a preventive architectural decision (open stack, API abstraction) can be made at negligible migration cost.**

◆ PHASE 2 · DEPENDENCY (YEAR 3) — Estimated lock-in score: 7/10 — CRITICAL INFLECTION POINT

Three-year Savings Plans and active Capacity Blocks create non-cancellable financial exposure. Data volumes have typically multiplied fivefold, rendering egress fees a six-figure item for mid-sized organisations. Between 8 and 12 proprietary services are interconnected, activating the superlinear technical lock-in effect. Teams hold multiple certification levels — passive resistance from middle management is now observable as a migration inhibitor. Migration remains plannable but costly (6–18 months, significant budget).

◆ PHASE 3 · CAPTIVITY (YEAR 5) — Estimated lock-in score: 9/10 — ECONOMIC POINT OF NO RETURN

Automatic contract renewal has incorporated terms unknown at the time of initial signing (Product Terms modified 60 times over 5 years for Azure). Egress costs exceed the annual IT budget of a mid-sized organisation. The entire architecture is built around proprietary managed services, making a full rewrite an 18–24 month project. Senior engineers have 5 years of exclusive platform experience — their probable departure in the event of a forced migration constitutes a documented HR risk. The cognitive layer is fully established: organisations can no longer define their requirements or measure their performance independently of the supplier's ecosystem.

4.2 The economic point of no return — formalisation

We define the economic point of no return (PNR) as the moment at which the total migration cost C_m exceeds the net present value of migration benefits B_m over the remaining commitment period. Formally: **PNR = {t | C_m(t) > B_m(t)}**, where C_m(t) = Σ(egress fees + residual commitments + refactoring + double-run + retraining) and B_m(t) = monthly post-migration savings × residual duration. Empirically, for organisations with more than 36 months of primary-hyperscaler tenure and active commitments, this ratio consistently exceeds 48 months — making migration economically unviable absent an external regulatory or strategic constraint.

This formalisation extends the "lock-in value" concept introduced by Shapiro & Varian (1999, op. cit.) to the specific structure of hyperscaler cloud contracts, where the asymmetry between entry cost (zero) and exit cost (positive and volume-dependent) systematically accelerates PNR as a function of data accumulation.

[11] Lock-in scores are indicative estimates constructed from the weighted aggregation of six layers. They are intended as a communication tool, not a scientifically validated measurement.

[12] The PNR formalisation is original. It operationalises Shapiro & Varian's (1999) lock-in value concept in the hyperscaler cloud contract context.

Amine RAITI · CC BY-NC-SA 4.0

5. COMPETITION LAW & COMPARATIVE REGULATORY RESPONSES

5.1 Market concentration and abuse of dominance

AWS, Azure and GCP jointly account for 67% of the global IaaS/PaaS cloud market (Synergy Research Group, Q1 2026). This concentration exceeds the collective dominance threshold established in EU case law (CJEU, *United Brands v Commission*, 1978[13]). The Digital Markets Act (DMA, Regulation (EU) 2022/1925) has designated both Microsoft and Alphabet as gatekeepers, though its provisions address market access rather than contractual retention mechanisms specifically. The EU Data Act (2023/2854) represents the most directly applicable instrument, specifically targeting obstacles to cloud switching in Articles 23–25.

In UK competition law, the Competition Act 1998 section 18 prohibits the abuse of a dominant position. The CMA's 2023 Cloud Market Study explicitly identified switching barriers as a primary competition concern, concluding that "the combination of technical barriers, egress fees and contract terms can make switching or multi-clouding excessively difficult"[14]. No formal enforcement action has followed from this finding at the date of this paper.

5.2 The EU Data Act as partial regulatory response

◆ WHAT THE DATA ACT COVERS — IN FORCE SINCE SEPT. 2025

ART. 23: accessibility and comprehensibility of contractual terms — directly applicable to linguistic substitution clauses. **ART. 25:** unenforceability of terms that create obstacles to switching cloud providers — applicable to non-cancellability clauses and disproportionate termination fees. **ART. 25(2):** data portability obligations — applicable to proprietary formats.

◆ IDENTIFIED GAPS — POLICY LACUNAE

LAYER 2 — PRICING: egress fees as a pricing-based switching obstacle are not explicitly covered — Art. 25 targets contractual clauses, not tariff structures. **LAYERS 4–6:** behavioural, decisional and cognitive obstacles fall outside the scope of the text. **COMPETENT AUTHORITIES:** no EU Member State had formally designated its Data Act competent authority as of 1 June 2026.

5.3 Comparative regulatory responses — FR · DE · UK

Regulator	Available instrument	Formal proceedings	Assessment
DGCCRF (FR)	Commercial Code L.442-1 · Consumer Code L.212-1	None public as of 01/06/26	Solid instrument · underused
Bundeskartellamt (DE)	GWB §19a · §19 · Active designations (Amazon 2022 · Microsoft 2023)	Active designations	Most powerful available instrument
CMA (UK)	Competition Act 1998 s.18 · Cloud Market Study 2023	Study published · no formal proceedings	Diagnosis made · follow-up insufficient

5.4 Cognitive lock-in as a cross-cutting regulatory gap

Layer 6 (cognitive lock-in) falls outside all existing regulatory frameworks. It involves no clause to challenge, no anticompetitive behaviour in the classical sense, and no obstacle to portability in the technical meaning of the Data Act. Yet it may represent the most durable barrier to competition in mature cloud markets — because it operates on organisations' capacity to conceive and evaluate alternatives. After five years of primary-hyperscaler dependency, an organisation's performance KPIs, benchmarking criteria and team evaluation frameworks are defined within the supplier's console. A migration implies not merely technical and financial costs but the loss of the organisation's own measurement standard. This observation suggests that an effective regulatory response to cloud lock-in must ultimately address organisational dependency beyond the contractual and technical dimensions currently covered by existing instruments. The UK case *Streetmap.EU Ltd v Google Inc* [2016][15] offers a precedent for competition law engaging with the cognitive dimensions of market dominance — though the context (search default settings) differs from the cloud dependency analysed here.

[13] CJEU, United Brands v Commission, 14 February 1978, Case 27/76, ECR 1978, p. 207.
[14] CMA, "Cloud Services Market Study — Final Report", January 2023, para. 4.58, available at gov.uk/cma.
[15] Streetmap.EU Ltd v Google Inc & Ors [2016] EWHC 253 (Ch), para. 112 (Roth J.).

Amine RAITI · CC BY-NC-SA 4.0

6. PROPOSALS — RESEARCH · POLICY · PRACTICE

6.1 Gaps in the existing literature

A review of the cloud vendor lock-in literature (Opara-Martins et al., 2016; Petcu, 2013; Armbrust et al., 2010; Farrell & Klemperer, 2007) confirms that Layers 1 to 3 of the proposed taxonomy are partially documented — primarily from a technical perspective and, to a lesser extent, an economic one. Layers 4 (Skills Lock-in), 5 (Demand Lock-in) and 6 (Cognitive Lock-in) are, to the author's knowledge, absent from the formalised academic literature on cloud computing. The linguistic substitution mechanism — the substitution of the client-language version by an English-only version declared legally binding — is not documented in any of the reviewed works. These gaps represent research opportunities of both academic and policy significance.

6.2 Research agenda

◆ EMPIRICAL RESEARCH

Empirical measurement of the economic point of no return by sector and organisation size. Longitudinal cohort study of skills evolution in organisations post-migration (3-year horizon). Comparative analysis of cloud contractual clauses over 10 years (2015–2025) to document temporal evolution of retention mechanisms. Measurement of the correlation between primary-hyperscaler tenure and capacity for functional requirement reformulation (test of Layer 6).

◆ LEGAL RESEARCH

Analysis of case law on linguistic substitution clauses in cross-border B2B contracts. Comparative study of regimes governing waiver of mandatory EU law rights by private contract. Qualification of cognitive lock-in (Layer 6) within existing competition law frameworks. Applicability of Data Act art. 25 to pricing-based switching obstacles (egress fees) — a question of direct practical significance not yet resolved by the regulatory text or its guidance documents.

6.3 Regulatory proposals

◆ FOUR PROPOSALS FOR REGULATORY FRAMEWORK EXTENSION — POLICY

Proposal 1 — Extension of Data Act art. 25 to pricing-based switching obstacles: the current wording of art. 25 targets contractual clauses that impede switching. An explicit extension to pricing structures creating a disproportionate economic obstacle (egress fees) — via delegated act or interpretative guidance — would cover Layer 2 of the taxonomy, currently outside scope. The Commission's forthcoming review of the Data Act implementation (expected 2027) provides a procedural vehicle for this extension.

Proposal 2 — Simultaneous multi-language publication obligation: require cloud service providers to publish their Terms of Service simultaneously in all EU official languages in which they actively market their services, with updates in non-English languages published within 5 business days of the English version, failing which the update is unenforceable in the relevant Member State(s). This measure addresses the structural temporal gap documented in Section 3.

Proposal 3 — Individual notification obligation for material modifications: require individual notification by documented electronic means for any material modification to Product Terms, with a minimum 60-day notice period and a right to terminate without penalty in response to a material modification. The absence of individual notification is one of the most significant gaps in the current cloud contractual framework.

Proposal 4 — Formal designation of Data Act competent authorities: no EU Member State had formally designated its Data Act competent authority as of 1 June 2026, eighteen months after the regulation's adoption and nine months after its applicability date. This implementation gap deprives organisations of any effective administrative remedy under the regulation and should be addressed as a matter of urgency by each Member State.

6.4 Conclusion

The six-layer taxonomy presented in this working paper maps the full architecture of retention mechanisms operating in hyperscaler cloud contracts. It reveals that the most durable mechanisms — Layers 5 and 6 — are precisely those which are least visible, least legally contestable, and least covered by existing regulatory frameworks. An effective regulatory response to cloud vendor lock-in cannot be complete without a theory of organisational digital dependency that integrates these behavioural and cognitive dimensions beyond the contractual and technical obstacles currently addressed by the Data Act, the DMA and national competition law instruments.

[16] *Regulation (EU) 2023/2854, art. 25(2): "Contractual terms [...] that prevent or restrict the right to switch to a different data processing service provider shall be deemed unenforceable."*

[17] *The Data Act applicability date (September 2025) is subsequent to its publication (OJ June 2024) and prior to the drafting of this paper (June 2026).*

Amine RAITI · CC BY-NC-SA 4.0

REFERENCES

Academic literature

Armbrust, Michael, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia. 2010. "A View of Cloud Computing." *Communications of the ACM* 53 (4): 50–58.

Barthélemy, Jérôme. 2003. "The Seven Deadly Sins of Outsourcing." *Academy of Management Executive* 17 (2): 87–100.

Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *Handbook of Industrial Organization*, vol. 3, edited by Mark Armstrong and Robert Porter, 1967–2072. Amsterdam: Elsevier.

Lacity, Mary C., and Leslie P. Willcocks. 1998. "An Empirical Investigation of Information Technology Sourcing Practices: Lessons from Experience." *MIS Quarterly* 22 (3): 363–408.

Opara-Martins, Justice, Reza Sahandi, and Feng Tian. 2016. "Critical Analysis of Vendor Lock-In and Its Impact on Cloud Computing Migration: A Business Perspective." *Journal of Cloud Computing: Advances, Systems and Applications* 5 (4).

Petcu, Dana. 2013. "Portability and Interoperability between Cloud Providers: Challenges and Case Study." In *Towards a Service-Based Internet*, Lecture Notes in Computer Science 8135. Berlin: Springer.

Porter, Michael E. 1980. *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. New York: Free Press.

Shapiro, Carl, and Hal R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press.

Legislation and regulatory instruments

European Parliament and Council. 2023. Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act). *OJ L* 143, 11 June 2024.

European Parliament and Council. 2022. Regulation (EU) 2022/1925 on contestable and fair markets in the digital sector (Digital Markets Act). *OJ L* 265, 12 October 2022.

European Parliament and Council. 2018. Directive (EU) 2018/1972 establishing the European Electronic Communications Code. *OJ L* 321, 17 December 2018.

France. 1994. Loi n°94-665 du 4 août 1994 relative à l'emploi de la langue française (Loi Toubon). *JORF* n°0180, 5 August 1994.

Germany. Gesetz gegen Wettbewerbsbeschränkungen (GWB), §§ 19, 19a. Version in force 1 June 2026.

United Kingdom. Unfair Contract Terms Act 1977 (UCTA), ss. 2–3. Consumer Rights Act 2015 (CRA), ss. 62–65. Competition Act 1998, s. 18.

Case law

CJEU. 1978. *United Brands v Commission*, Case 27/76, ECR 1978, p. 207. 14 February 1978.

UK High Court (Ch). 2016. *Streetmap.EU Ltd v Google Inc & Ors* [2016] EWHC 253 (Ch). Roth J.

Institutional reports

Bundeskartellamt. 2022. Proceeding against Amazon pursuant to §19a GWB. Decision of 5 July 2022. bundeskartellamt.de.

Bundeskartellamt. 2023. Proceeding against Microsoft pursuant to §19a GWB. Decision of 12 December 2023. bundeskartellamt.de.

Competition and Markets Authority. 2023. *Cloud Services Market Study — Final Report*. January 2023. gov.uk/cma.

Synergy Research Group. 2026. Cloud Market Share Q1 2026. Published April 2026.

Primary contractual documents

47 contractual documents (General Terms, Product Terms, DPAs, SLAs, Addenda) published by AWS, Google Cloud and Microsoft Azure in FR/EN/DE, collected between January 2022 and June 2026. Full corpus in Annex A. Sources: aws.amazon.com/legal/ · cloud.google.com/terms · microsoft.com/licensing

Amine RAITI · CC BY-NC-SA 4.0

ANNEX A — PRIMARY DOCUMENTARY CORPUS

Public documents collected between January 2022 and June 2026 · Official sources of the three hyperscalers

A.1 — Amazon Web Services (16 documents)

Document	Versions analysed	Languages	Critical clauses	Layer
AWS Customer Agreement (ToS)	15 Apr. 2026 · 22 May 2026 · 1 Jun. 2026	FR · EN · DE	s.1.4 · s.5.4.2 · s.5.5	1 · 2
AWS Product Terms	Jan. 2022 → Jun. 2026	EN (only)	s.50.12.1–3 · s.50.14 · s.50.15 · s.50.17	1 · 2
AWS Data Act Addendum	Single version 2024	EN · FR	s.1.28 · portability obligations	1
AWS UK Switching Addendum	Single version 2024	EN	s.1.30 · switching facilitation	1 · 2
AWS Data Processing Addendum	2023 · 2024 · 2025	FR · EN · DE	Bedrock AI services opt-out	1
AWS SLAs (Service Level Agreements)	EC2 · S3 · RDS · EKS · Lambda	EN (only)	Liability exclusions · Force majeure	1

A.2 — Google Cloud Platform (15 documents)

Document	Versions analysed	Languages	Critical clauses	Layer
GCP Terms of Service	Aug. 2023 (FR) · Jun. 2026 (EN)	FR · EN · DE	s.14.18 · s.8.8 · s.2.6	1 · 2
GCP Service Specific Terms	2022 → 2026 (13 EN updates missing from FR)	EN (FR frozen)	BigQuery · Vertex AI · Spanner pricing	1 · 2 · 3
GCP Data Processing Addendum	2023 · 2025	FR · EN · DE	AI Opt-out · GDPR art. 28	1
GCP SLAs (8 services)	GKE · BigQuery · Cloud SQL · Cloud Run · GCS · Pub/Sub · Vertex · Compute	EN (only)	Service credit exclusions · Downtime definition	1
GCP Committed Use Contracts	1-year · 3-year · 2023–2026	EN · FR	s.8.8 · no refund on termination	1 · 2

A.3 — Microsoft Azure (16 documents)

Document	Versions analysed	Languages	Critical clauses	Layer
Microsoft Online Services Agreement (MOSA)	2019 (signed) + monthly amendments 2019–2026	FR · EN	Washington State governing law · Capacity Blocks	1
Product Terms MOSA	Jan. · Feb. · Apr. · Jun. 2026	FR · EN	608 modifications · Capacity Blocks Jun. 2026	1
Enrollment for Education Solutions (EES)	Jan. 2026 (EN/FR) · Jun. 2026 (EN/FR)	FR · EN	Waiver Dir. 2018/1972 · absent from FR version	1
Azure Savings Plans	2023 · 2024 · 2025	FR · EN	Non-cancellable · 12% early termination fee	1 · 2
Azure DPA · SLA · Offer Details	AKS · Azure SQL · Blob · OpenAI · Monitor	EN (only)	Liability exclusions · CLOUD Act exposure	1 · 3

Amine RAITI · CC BY-NC-SA 4.0

ANNEX B — LINGUISTIC SUBSTITUTION: PRIMARY DATA

B.1 — AWS Customer Agreement: 22 May 2026 vs 1 June 2026 — comparative analysis

This comparison is of particular relevance as it documents the only substantive legal modification between two consecutive versions — a modification that would have gone undetected without systematic monitoring.

Section	Content	22 May 2026	1 June 2026	Modification	Layer
Header	Linguistic substitution clause	Present · EN prevails	Present · EN prevails	**Identical**	1
s.1.4	Suspension without notice	Present	Present	**Identical**	1
s.1.24	Bedrock AI services list	"Amazon Quick Suite (incl. QuickSight)"	"Amazon Quick"	**Label simplified**	1
s.5.4.2	Savings Plans noncancellable	Present	Present	**Identical**	1 · 2
s.5.5	Capacity Blocks nonrefundable	Present	Present	**Identical**	1 · 2
s.50.12.2	Bedrock abuse detection	"automated abuse detection mechanisms"	"store inputs/outputs 30 days" + "human review" option	**MODIFIED — substance**	1
s.50.17	Amazon Quick	Absent	Present · capacity pricing mandatory for automation	**NEW SECTION**	1 · 2
s.1.28	Data Act Addendum	Present	Present	**Identical**	1
s.1.30	UK Switching Addendum	Present	Present	**Identical**	1 · 2

B.2 — GCP Terms of Service: FR vs EN translation delays (2022–2026)

Language	Last substantive update	Equivalent EN version	Delay	Missing EN updates
French (FR)	August 2023	June 2026	**2 years 9 months**	**13 updates**
German (DE)	June 2022	June 2026	**4 years**	**16 updates**

B.3 — Azure EES: translation delays self-admitted by Microsoft (2026)

Version	Published EN	Published FR	Delay	Self-admitted source
EES Jan. 2026	1 February 2026	9 February 2026	**8 days**	EES document: "Published February 1, 2026 in English, translated February 9, 2026"
EES Jun. 2026	1 June 2026	~9 June 2026 (estimated)	**~8 days**	Admitted in Jun. 2026 EES document
Dir. 2018/1972 waiver	Present EN Jan. + Jun. 2026	**Absent FR Jan. + Jun. 2026**	**Not translated**	**Clause not present in FR version at either date**

Note: Publication dates are extracted from metadata and internal mentions within the contractual documents themselves. The algorithmic diff methodology used to compare versions is described in Annex D.

Amine RAITI · CC BY-NC-SA 4.0

ANNEX C — TWELVE-QUESTION LOCK-IN DIAGNOSTIC

Organisational communication tool · Not empirically validated at this stage · Proposed as a basis for future research

Q.	Layer	Diagnostic question · Yes = 2 pts · No = 0 pt	If yes	Refer to
1a	**Contractual**	Do you have Savings Plans or Capacity Blocks expiring beyond 12 months?	**2 pts**	Guide A · Legal
1b	**Contractual**	Does the non-English version of your cloud ToS state that the English version prevails?	**2 pts**	Guide A · Legal
2a	**Pricing**	Have you calculated the egress cost of moving all your current data out of your primary cloud provider?	**2 pts**	Guide C · CFO
2b	**Pricing**	Do you know the early termination fee applicable to your primary cloud commitments?	**2 pts**	Guide C · CFO
3a	**Technical**	Do more than 5 services in your architecture use APIs without an S3-compatible equivalent?	**2 pts**	Guide B · CTO
3b	**Technical**	Do you have a documented dependency map showing which services use proprietary APIs?	**2 pts**	Guide B · CTO
4a	**Skills**	Do more than 50% of your cloud engineers hold a proprietary certification without an open equivalent?	**2 pts**	Guide E · HR
4b	**Skills**	Have you identified the resistance profiles within your technical teams?	**2 pts**	Guide E · HR
5a	**Demand**	Did the last 3 business requests to your IT department name a specific product (Azure OpenAI, BigQuery…)?	**2 pts**	Guide E · PO
5b	**Demand**	Do your Product Owners systematically apply the 5 sovereignty questions before validating a stack?	**2 pts**	Guide E · PO
6a	**Cognitive**	Are your IT performance KPIs defined and measured within your primary cloud provider's console?	**2 pts**	Guide D · All
6b	**Cognitive**	Would you be unable to measure IT performance during a 6-month migration?	**2 pts**	Guide D · All

0–6 pts	7–12 pts	13–18 pts	19–24 pts
Low lock-in · Preventive action · Guide B	Moderate · Guides A+B+C recommended	High · Guide D · Plan migration	Structural · All guides · External support

Methodological note: equal weighting of 12 questions (2 points each) is a simplification. Future research could develop differentiated weighting based on empirical measurement of actual migration costs per layer.

Amine RAITI · CC BY-NC-SA 4.0

ANNEX D — METHODOLOGY OF CONTRACTUAL DOCUMENT COMPARISON

D.1 Document collection

Contractual documents were collected directly from the official websites of the three providers (aws.amazon.com/legal/ · cloud.google.com/terms · microsoft.com/licensing) between January 2022 and June 2026. For each provider, versions available in the three languages (FR · EN · DE) were downloaded upon each detected update, via systematic monitoring of the terms pages. Publication dates are extracted either from document metadata or from explicit internal mentions (e.g. Azure EES: "Published February 1, 2026 in English, translated February 9, 2026").

D.2 Document pre-processing

HTML documents were pre-processed in two stages: (a) text extraction by removal of HTML tags, scripts and CSS styles; (b) text normalisation by replacement of HTML entities, non-breaking spaces and typographic ligatures (e.g. ﬁ → fi), and compression of multiple whitespace characters. This normalisation step is essential to distinguish substantive modifications from formatting changes — which represent the majority of "modifications" detected by a raw HTML diff. All differences identified as substantive were independently confirmed in the original documents to exclude normalisation artefacts.

D.3 Algorithmic comparison

Comparison between versions was conducted using normalised textual diff, section by section (each section identified by its number: s.1.4, s.5.5, s.50.12.2, etc.). For each pair of sections (version X vs version Y), normalised Levenshtein distance was computed to flag substantive modifications. Section pairs exhibiting greater than 95% similarity after normalisation were classified as "identical in substance" — residual typographic and formatting artefacts not constituting legally significant modifications.

D.4 Manual legal review

All modifications identified as substantive by the algorithm were subject to manual review to: (a) confirm that the difference constitutes a legally significant modification and not a normalisation artefact; (b) characterise the nature of the modification (addition · deletion · reformulation · change in scope); (c) identify the relevant lock-in layer under the six-layer taxonomy. The results of this manual review are reproduced in Annex B for linguistic comparisons and in the corpus of 19 individual legal analyses of which this study constitutes the synthesis.

D.5 Identified methodological limitations

◆ THREE LIMITATIONS EXPLICITLY ACKNOWLEDGED

Limitation 1 — Documentary instability: Azure's monthly modification of Product Terms (608 occurrences over 6 months) creates a moving corpus that is difficult to stabilise for longitudinal analysis. Conclusions on Azure modifications refer to the snapshots collected at the stated dates and may not reflect the state of the document on an intermediate date.

Limitation 2 — Non-public documents: Enterprise Agreements and individually negotiated contracts, which may contain significantly different clauses from the public standard terms, are excluded from the analysis. The study covers exclusively conditions applicable to the majority of non-negotiating clients.

Limitation 3 — Legal qualification: the legal characterisation of analysed clauses (voidness, significant imbalance, violation of mandatory law) rests on the author's independent analysis and not on a court decision or formal legal opinion. These qualifications should be verified by a qualified legal practitioner before any action is taken.

This working paper is based on the independent analysis of publicly available contractual documents.
It does not constitute legal advice. The author has no affiliation with any of the companies or institutions mentioned.

Nemo supra legem est · Summum ius, summa iniuria — Cicero

Public document · CC BY-NC-SA 4.0

AI Powered by Amine — AI is an amplifier of ideas and forms, not the source.

MAPPING THE STRATAGEMS OF VENDOR LOCK-IN

AWS · Google Cloud · Microsoft Azure

◆ **DOCUMENTED FRENCH CASE STUDY**

This study documents a real case observed on the French market, under French/EU law. Figures are in euros; legal references are to French and EU statutes. The mechanism it illustrates is recognised as universal across jurisdictions — the specific numbers, legal citations, and institutional frameworks (CSE, OPCO, RNCP, ANSSI, DGCCRF, etc.) are French-specific and are preserved here as documentary evidence, not as an action guide for another jurisdiction.

◆ WHAT THIS TOOL IS

Internal diagnostic tool for General Management, IT Management, Legal, Finance, HR and Product Owners. Six layers · three hyperscalers · five blindness profiles · a five-year temporal dimension. Based on comparative analysis of 47 public contractual documents (2022–2026). This document is not legal advice.

◆ SUMMARY

Page 2 — the 6 layers × 3 hyperscalers × intensity. Page 3 — the temporal dimension: years 1, 3, 5. Page 4 — mapping invisibility by profile. Page 5 — diagnostic grid and action plan.

STRATAGEM

Amine RAITI · CC BY-NC-SA 4.0

THE 6 LAYERS × 3 HYPERSCALERS × INTENSITY × VISIBILITY

◆ READING RULE — THE STRATIGRAPHIC MAP

Like a geological cross-section: surface layers (1–2) are visible provided the right documents are read. Middle layers (3–4) are visible only to specific profiles with the right tools. Deep layers (5–6) are invisible to everyone without a dedicated analytical framework. Lock-in deepens over time: the longer an organisation uses the service, the denser and harder to cross each layer becomes. The economic point of no return is generally between Year 3 and Year 4.

Layer · Mechanism	Description	AWS	GCP	Azure	Visible by · when
1 · Contractual	Language clause · Non-cancellability · Unilateral modification	●●● s.5.4.2·5.5·50.12.3	●●●● s.14.18·s.8.8	●●●● MOSA·EES·6 layers	Legal — if reading the English version and the Product Terms
2 · Pricing	Egress fees · Termination charges · Entry/exit asymmetry	●●● $0.09/GB · NAT €6.7k	●●● $0.12/GB · BQ $5/TB	●●●● 12% termination + $50k/yr	Finance — if the exit cost is quantified before signing
3 · Technical	Proprietary APIs · Managed services · Superlinear effect	●●● Lambda·DynamoDB·Bedrock	●●●● BigQuery·Vertex·Spanner	●●● Cosmos·Entra·Synapse	IT — if API dependencies are audited (rarely done)
4 · Skills	Over-certification · Attachment bias · Passive resistance	●●● 12 AWS certifications	●● 11 GCP certifications	●●● 9+ Azure certifications	HR — if team certifications are mapped
5 · Demand	Requests for named products · Prescribing business units	●● Named Bedrock·EKS	●●● Named BigQuery·Looker	●●●● M365 = entire ecosystem	PO/Business — if needs are reformulated functionally
6 · Cognitive	Proprietary KPIs · Loss of measurement baseline	●●● CloudWatch·Cost Explorer	●● Cloud Monitoring	●●●● Excel·Teams·AD	20 years 🚫 Nobody — invisible by design, the most irreversible

◆ CENTRAL PARADOX — INVISIBILITY AT THE RIGHT MOMENT

Each lock-in layer is designed to be invisible to the profile that would have the power to stop it at the right moment. The lawyer does not read the monthly Product Terms. Finance does not quantify egress fees before signing. HR does not map certifications before migration. The PO does not reformulate needs functionally. General management receives summaries that aggregate without analysing. Lock-in is built in organisational blind spots.

STRATAGEM

Amine RAITI · CC BY-NC-SA 4.0

THE TEMPORAL DIMENSION — HOW LOCK-IN DEEPENS

◆ YEAR 1 · ADOPTION — 2/10

Layer 1: short commitments, low contractual exposure. Layer 2: low data volume, no significant egress. Layer 3: 2–3 proprietary services, migration possible in a few weeks. Layer 4: certifications underway, no strong bias yet. Layer 5: first named requests, still few precedents. Layer 6: nascent KPIs, little history, comparison still possible. ✅ Migration possible within weeks · Fiche B is sufficient.

◆ YEAR 3 · DEPENDENCE — 7/10

Layer 1: 3-year Savings Plans, active Capacity Blocks, non-cancellable commitment. Layer 2: volume ×5, six-figure egress for a large organisation. Layer 3: 8–12 interlocking services, superlinear effect activated. Layer 4: teams certified at 3 levels, passive management resistance. Layer 5: 2–3 years of precedent, established pattern in briefs. Layer 6: 3 years of KPI history, loss of external comparison. ⚠ Migration: 6–18 months · Fiches A+B+C+D required.

◆ YEAR 5 · CAPTIVITY — 9/10

Layer 1: auto-renewal, Product Terms modified 60 times, unknown clauses. Layer 2: prohibitive egress, data-migration cost exceeds an SME's annual IT budget. Layer 3: entire architecture proprietary, rewrite = 18–24 months. Layer 4: senior engineers 5 years on the platform, likely departure if migration is forced. Layer 5: business units no longer conceive needs outside the ecosystem. Layer 6: impossible to compare pre/post performance — baseline lost. ❌ Migration: 18–24 months · organisational risk · external support recommended.

◆ THE ECONOMIC POINT OF NO RETURN — DEFINITION

The economic point of no return is the moment the total migration cost (egress fees + remaining commitments + refactoring + double-run + training) exceeds the expected benefit (future savings × remaining commitment period). Formula: migration break-even = (egress cost + remaining commitments + refactoring + double-run) ÷ (post-migration monthly savings). If this ratio exceeds 36 months, migration is only worthwhile if an external regulatory or strategic constraint forces it. For most organisations in Year 5, this ratio exceeds 48 months.

◆ THE 18-MONTH RULE

The dependency audit (Fiche B) should be conducted every 18 months. Between Year 1 and Year 3, the action window is open and the cost of a partial migration remains manageable. Past Year 3, every additional month on the platform reduces the economic viability of exit. The decision to migrate must be taken 2–3 times earlier than the organisation thinks necessary — because migration time is systematically underestimated.

STRATAGEM

Amine RAITI · CC BY-NC-SA 4.0

MAPPING INVISIBILITY — BY PROFILE, BY LAYER, BY MOMENT

Profile	What they see	What they don't see	Possible moment of action	When they discover it
⚖ Legal	Visible contractual clauses · Layer 1 (if reading English version)	Monthly Product Terms · Layers 4/5/6 · Unilateral modification	Before signing · Never after	At first dispute — too late
🖥 IT	Layers 2·3 · APIs · Technical dependencies (if audited)	Fine layer 1 · Layers 4·5·6 · Temporal dimension of lock-in	At architectural choice · Year 1	When migration is requested · Year 3–5
💼 Finance	Monthly invoice · Partial layer 2	Exit cost · Future egress fees · Termination charges · Double-run · Layers 3·4·5·6	At signing · Never after	When migration budget is presented — shock
👥 HR	Tech turnover · Layer 4 (if certifications mapped)	Skills lock-in as cause of turnover · Passive middle-management resistance · Layers 1·2·3·5·6	During certification mapping · Years 1–2	When the migration project fails — too late
🎯 PO/Business	Features · UX · Time-to-market	That they generate Layer 5 · Contractual commitment created by their requests · Layers 1·2·3·4·6	Before each new business request	Never spontaneously — only if Fiche E is shared
★ General Management	IT summary report · Total monthly cost	None of the 6 layers in full — this document is designed for them	At all times, if the full file is presented	During an acquisition · financial crisis · regulatory injunction

◆ THE PARADOX OF THE WELL-RUN ORGANISATION

An organisation where every profile does their job correctly can still build structural lock-in. The lawyer reads the clauses submitted to them. Finance controls the invoice they are shown. IT manages the architecture entrusted to them. HR manages the teams presented to them. The PO expresses the needs they are asked to express. None of them does their job badly — but lock-in is built in the spaces between their responsibilities. That is why only a cross-cutting reading, like this document, can reveal its full architecture.

STRATAGEM

Amine RAITI · CC BY-NC-SA 4.0

INTERNAL DIAGNOSTIC GRID — 12 QUESTIONS · SCORE OUT OF 24

◆ SCORING

Yes = 2 pts · No = 0 pt

#	Diagnostic question	If yes →
1a	Do you have Savings Plans or Reserved Instances expiring in more than 12 months?	→ Fiche A · Legal
1b	Does the French version of your cloud terms state the English version prevails?	→ Fiche A · Legal
2a	Have you calculated the egress cost of exporting all your current data?	→ Fiche C · Finance
2b	Do you know the early-termination fees on your main cloud commitments?	→ Fiche C · Finance
3a	Do more than 5 services in your architecture use APIs with no S3-compatible equivalent?	→ Fiche B · IT
3b	Do you have a documented dependency grid by cloud service used?	→ Fiche B · IT
4a	Do more than 50% of your cloud engineers hold a proprietary certification with no open equivalent?	→ Fiche E · HR
4b	Have you identified the 5 resistance profiles in your technical teams?	→ Fiche E · HR
5a	Did the last 3 business requests to IT mention a named product (Azure OpenAI, BigQuery…)?	→ Fiche E · PO
5b	Do your POs systematically ask the 5 sovereignty questions before validating a stack?	→ Fiche E · PO
6a	Are your IT performance KPIs defined and measured in your main provider's console?	→ Fiche D · All
6b	Would you be unable to measure IT performance during a 6-month migration?	→ Fiche D · All

◆ SCORE INTERPRETATION

0–6 pts: weak lock-in — Fiche B is sufficient, act preventively. 7–12 pts: moderate lock-in — Fiches A+B+C recommended, audit within 3 months. 13–18 pts: high lock-in — Fiche D required, plan migration. 19–24 pts: structural lock-in — all fiches, external support recommended.

◆ SOURCE

Based on 47 public contractual documents from AWS · GCP · Azure (2022–2026). Summum ius, summa iniuria — Cicero.

STRATAGEM

Amine RAITI · CC BY-NC-SA 4.0

🇬🇧 ENGLISH VERSION · RECIPIENT: CMA · FCA · ICO

THE CONTRACTUAL WAR MACHINE

Complete Anatomy · English Version

Systemic deconstruction of hyperscaler vendor lock-in
based on nineteen legal analyses of AWS · GCP · Azure terms
addressed to the CMA · FCA · ICO · European regulators

UK LEGAL BASIS

- ◆ UCTA 1977 s. 3 (reasonableness test)
- ◆ SGSA 1982 s. 13 (implied terms)
- ◆ Hadley v Baxendale [1854]
- ◆ Smith v Eric S Bush [1990]
- ◆ EU Data Act (2023/2854) art. 23 · 25
- ◆ CRA 2015 s. 62 (unfair terms)
- ◆ Competition Act 1998 s. 18
- ◆ Interfoto v Stiletto [1989]
- ◆ CMA Cloud Market Study 2023
- ◆ UK GDPR art. 5 · 6 · 13-14

AMINE RAITI · Infrastructure Architect & SRE · Paris

Public document · CC BY-NC-SA 4.0 · Not legal advice

MACHINA

DESIGN PRINCIPLE · THREE-LAYER ARCHITECTURE · CAPTIVITY CYCLE

⚠ CORE FINDING

These contracts are not poorly drafted. They are **PERFECTLY DRAFTED** for a precise objective that is not the customer's: to **MAXIMISE EXIT COSTS** while minimising entry costs. This is not a side effect — it is the primary function. Every clause analysed becomes immediately explicable when read with this objective in mind rather than through the lens of a good-faith buyer.

01 CONTRACTUAL	02 PRICING	03 TECHNICAL
Non-cancellability · Language clause · Unilateral modification · Waiver of legal rights	Egress fees 0.08–0.12 $/GB · Entry/exit asymmetry · Early termination fees 12% (Azure)	Proprietary APIs · Managed services · Superlinear portfolio dependency

◆ INTERACTION PRINCIPLE

A customer who negotiates a cancellable contract remains locked in by egress fees. A customer who circumvents egress fees remains locked in by non-cancellable commitments. A customer who builds an open architecture remains locked in by monthly Product Terms updates. Exit is systematically costly regardless of strategy.

CAPTIVITY CYCLE — 5 PHASES

ENTERS: Free credits · Excellent APIs · Zero entry cost → **SETTLES**: Good services · Perfect docs → **DEPENDS**: APIs in codebase · Teams certified → **COMMITS**: −30% over 3 yrs · Non-cancellable → **DISCOVERS**: Too late · Three layers active

◆ THE RATIONAL CHOICE THAT BECAME A TRAP — STRUCTURAL

Reserved Instances cost 30–40% less than on-demand. The saving is real. It comes at the price of non-cancellability. The customer trades freedom for a price reduction — a rational short-term choice that becomes a medium-term trap. **UCTA 1977 s. 3**: unreasonable exclusion of liability = unenforceable. **CRA 2015 Sch. 2**: terms creating significant imbalance = unfair. **EU Data Act art. 25**: contractual obstacles to switching cloud providers = unenforceable.

Amine RAITI · CC BY-NC-SA 4.0

LANGUAGE · CASCADE INCORPORATION · NON-CANCELLABILITY · PRICE ASYMMETRY

◆ MEC. I · LANGUAGE CLAUSE — 3 STRATEGIES, 1 RESULT

CRITICAL · ALL HYPERSCALERS

AWS (HEADER OF EVERY PAGE — FR VERSION):

> *"The following translation is provided for informational purposes only. [...] the English version will prevail." — Gap: 37 days · Explicit · Unlawful.*

GCP (SECTION 14.18, FROZEN AUGUST 2023):

> *"[...] the English text will govern." — FR gap: 2 yrs 9 mths · 13 missed updates · Concealed.*

AZURE MOSA (6-LAYER FRAGMENTATION, MONTHLY PRODUCT TERMS):

> *"Published February 1, 2026 in English, translated February 9, 2026." — Self-admitted · Documentary dissolution.*

◆ LEGAL BASIS

INTERFOTO V STILETTO [1989]: onerous terms must be drawn to attention · **CRA 2015 SCH. 2 PARA. 1(I)** · **EU DATA ACT ART. 23**

◆ MEC. II · CASCADE INCORPORATION — CLAUSE OUTSIDE SIGNED DOCUMENT

HIGH

EES MICROSOFT · "ABOUT THIS DOCUMENT" (EN VERSION ONLY):

> *"Updates that Microsoft makes from time to time to Use Rights apply to Customer as set forth in Customer's agreement." — Absent from French version. Enforceable via incorporation by reference.*

◆ LEGAL BASIS

INTERFOTO V STILETTO [1989] · **UCTA 1977 S. 3**: customer must be given fair notice of terms · **CRA 2015 S. 68**: terms must be transparent and prominent

◆ MEC. III · SELECTIVE NON-CANCELLABILITY — THE CORE LOCK-IN

CRITICAL · AWS + AZURE

AWS SECTION 5.5 · CAPACITY BLOCKS:

> *"cannot be canceled nor modified [...] full price nonrefundable [...] last 30 minutes we may terminate your instances without notice."*

AZURE EES · SAVINGS PLAN (IDENTICAL FR + EN):

> *"noncancellable. Customer will owe the amount charged [...] even if the Agreement is terminated."*

◆ LEGAL BASIS

UCTA 1977 S. 3: reasonableness test · **HADLEY V BAXENDALE [1854]** · **CRA 2015 SCH. 2** · **EU DATA ACT ART. 25**

◆ MEC. IV · PRICE ASYMMETRY — SUPPLIER MODIFIES, CUSTOMER LOCKED

CRITICAL

AWS SECTION 50.12.3 · BEDROCK:

> *"modify the pricing [...] at any time. [...] charged for the duration [...] even if you terminate this Agreement."*

GCP SECTION 2.6 (ABSENT FROM FR VERSION 2023):

> *"Google may change the Fees at any time." — Unknown to French customers but enforceable.*

◆ LEGAL BASIS

UCTA 1977 S. 3 · **CRA 2015 S. 62** · **SMITH V ERIC S BUSH [1990]** · **CA 1998 S. 18** (abuse of dominant position)

Amine RAITI · CC BY-NC-SA 4.0

SUSPENSION · RIGHTS WAIVER · RATCHET EFFECT · EGRESS FEES

◆ MEC. V · SUSPENSION WITHOUT NOTICE — HIGH

AWS Section 1.4:

> *"remove or disable access [...] without prior notice."*

The threat of suspension is sufficient to deter any challenge. Azure MOSA Section 3.c provides 30 days' notice for non-payment — the one point where Azure outperforms AWS on this criterion. **CA 1998 s. 18 · CRA 2015 Sch. 2 para. 1(b)**: right to suspend must not be exercised in a disproportionate manner.

◆ MEC. VI · WAIVER OF EU RIGHTS — ENGLISH VERSION ONLY — CRITICAL · EES MICROSOFT

EES Microsoft — English version only (Jan + Jun 2026):

> *"Customer agrees to waive any and all entitlements [...] under the European Electronic Communications Code (Directive 2018/1972) Article 102 paras 1, 3, 5 ; Article 105 para 1 ; Article 107 paras 1 and 3."*

◆ LEGAL BASIS

The French Ministry of Education waives EU legal protections without this waiver being presented to it in its own language. **Interfoto v Stiletto [1989]**: onerous clauses must be explicitly drawn to the other party's attention before or at the time of contracting.

◆ MEC. VII · RATCHET EFFECT — CONTRACT WRITTEN AFTER SIGNING — CRITICAL · AZURE

EES Microsoft · Summary of Changes · June 1, 2026:

> *"New terms added for Azure Capacity Blocks." — Absent from Jan/Feb/Apr 2026 versions. Applied automatically to all existing contracts.*

◆ LEGAL BASIS

UCTA 1977 s. 3 · Interfoto v Stiletto [1989] · CMA Cloud Market Study 2023: identified switching barriers as primary competition concern.

GRADUS IV · EGRESS FEES — THE WALL OUTSIDE THE CONTRACT

Provider	Egress/GB	100 TB	In T&Cs?
AWS	$0.08–0.09/GB	**£7–8k**	**NO**
GCP	$0.08–0.12/GB	**£7–10k**	**NO**
Azure	$0.07–0.09/GB	**£6–8k**	**PARTIAL**
1 Petabyte	**£70,000–90,000 before technical migration**	**Not contractualised**	

Entry is free, exit is priced. The machine deliberately separates contractual from pricing to maximise opacity at signature. **EU Data Act art. 25 · CA 1998 s. 18**: obstacle to switching = potentially unenforceable.

Amine RAITI · CC BY-NC-SA 4.0

◆ GRADUS V · TECHNICAL LAYER

Proprietary APIs: Lambda/Cloud Functions/Azure Functions = non-standardised. Migration requires full rewrite.
Managed services: BigQuery/Redshift/Azure Synapse = proprietary SQL dialects. 10 years of queries = months of migration.
Superlinear effect: each additional service increases the cost of exiting all others.

◆ GRADUS VI · 4 REASONS THIS IS DIFFERENT

I — Scale: 65%+ global market → potential abuse of dominance. **CA 1998 s. 18.**
II — Information asymmetry: egress fees outside T&Cs · Directive 2018/1972 waiver in English only · constitutive, not accidental.
III — Positive law violated: UCTA 1977 · CRA 2015 · CA 1998 s.18 · EU Data Act art. 23–25.
IV — Deliberate design: drafted by top-tier legal teams · tested across jurisdictions · the waiver clause is not in English only by accident. CMA 2023 Study identified switching barriers — no enforcement followed.

Mechanism	AWS	GCP	Azure
Language substitution	Explicit · every page	Concealed s.14.18	Fragmented · 6 layers
Non-cancellability	s.5.4.2 · 5.5 · 50.12.3	s.2.1 · commitments	Savings Plans · Blocks
Refund on termination	Pro-rata if AWS terminates ✓	None · s.8.8	12% fee + $50k/yr cap
Data Act Addendum	✓ s.1.28	Absent	Absent MOSA/EES
EU rights waiver	Partial	Partial	Dir. 2018/1972 · EN only

◆ GRADUS VII · CONCLUSION — CLAUSE VS SYSTEM

An unfair clause is a problem for the civil courts. A system deliberately designed to create dependency, relying on unlawful clauses, at the scale of a 65% concentrated market — that is a problem for the competition regulator. The **CMA** holds the **Competition Act 1998 s. 18**. Its own 2023 Cloud Market Study identified switching barriers as a primary concern. No formal enforcement has followed. This document and the nineteen analyses underpinning it constitute the two faces of the same case: clause by clause, and as a coherent system. Both are necessary for regulators to act.

Nemo supra legem est · Summum ius, summa iniuria — Cicero

Public document · CC BY-NC-SA 4.0 · Not legal advice · Amine RAITI

SYSTEMIC STUDY · PUBLIC DOCUMENT · OPERATION DINDON

THE CONTRACTUAL WAR MACHINE

Complete Anatomy

◆ DOCUMENTED FRENCH CASE STUDY

This study documents a real case observed on the French market, under French/EU law. Figures are in euros; legal references are to French and EU statutes. The mechanism it illustrates is recognised as universal across jurisdictions — the specific numbers, legal citations, and institutional frameworks (CSE, OPCO, RNCP, ANSSI, DGCCRF, etc.) are French-specific and are preserved here as documentary evidence, not as an action guide for another jurisdiction.

◆ WHAT THIS STUDY IS

Systemic deconstruction of hyperscaler vendor lock-in, based on comparative analysis of nineteen legal studies of AWS, Google Cloud and Microsoft Azure terms of service. 3 lock-in layers · 7 documented mechanisms · 19 source legal analyses · 65% combined global market share.

◆ SOURCES

AWS ToS FR vs EN (15 Apr / 22 May 2026) · AWS ToS DE vs EN (German law) · AWS ToS EN vs British law · GCP ToS FR vs EN (Aug 2023 / June 2026) · GCP ToS DE vs EN (4-year lag) · GCP ToS EN vs British law · Azure MOSA — dispersion architecture · Product Terms EES FR vs EN (Jan/June 2026).

MACHINA

Amine RAITI · CC BY-NC-SA 4.0

GRADUS I — THE DESIGN PRINCIPLE

◆ FUNDAMENTAL FINDING

These contracts are not badly written. They are perfectly written — for a purpose that is not the customer's. An ordinary commercial contract frames a mutually beneficial exchange relationship. These contracts exist to maximise exit cost while minimising entry cost. This is not a side effect. It is the primary function. Every clause analysed makes immediate sense once the contract is read with that objective in mind, rather than with the assumption of a good-faith buyer.

◆ GRADUS II — THE THREE SUPERIMPOSED, CUMULATIVE LAYERS

The machine is designed so that failure at one layer is compensated by the other two. 01 Contractual: non-cancellability · language clause · unilateral modification · waiver of legal rights. 02 Pricing: egress fees $0.08–0.12/GB · entry/exit asymmetry · 12% early-termination fee. 03 Technical: non-portable proprietary APIs · managed services · superlinear dependency.

◆ LAYER INTERACTION PRINCIPLE

A customer who negotiates a cancellable contract remains locked in by egress fees and proprietary APIs. A customer who avoids egress fees remains locked in by non-cancellable commitments. A customer who builds an open architecture remains locked in by Product Terms modified monthly. Exit is systematically costly, whatever strategy is adopted.

◆ THE CAPTIVITY CYCLE — 5 PHASES

ENTERS FREELY — low entry costs, free credits, excellent APIs → SETTLES IN — good services, perfect documentation → DEPENDS — APIs in the code, certified teams → COMMITS — −30% over 3 years, non-cancellability → DISCOVERS — too late, three layers simultaneously.

◆ THE RATIONAL CHOICE TURNED TRAP

Reserved Instances cost 30–40% less than on-demand. The saving is real. It is obtained at the price of non-cancellability. The customer trades freedom for a price discount — a rational short-term choice that becomes a medium-term trap. The machine does not force — it convinces.

MACHINA

Amine RAITI · CC BY-NC-SA 4.0

MECHANISMS I TO IV — THE CONTRACTUAL LAYER

◆ MECHANISM I — THE LANGUAGE CLAUSE: THREE STRATEGIES, ONE RESULT

🇺🇸 AWS — header on every page (FR, 15 Apr 2026): "The translation below is provided for informational purposes only [...] the English version will prevail." Strategy: explicit · lag: 37 days · illegal, transparent. 🔵 GCP — Section 14.18, hidden (FR frozen Aug 2023, EN June 2026): "If this Agreement is translated [...] the English version prevails." Strategy: concealed · FR lag: 2y9m · DE lag: 4 years · illegal, opaque. 🔷 Azure — fragmented into 6 layers (MOSA + Product Terms + DPA + SLA + Offer Details + Preview): "These terms were published on February 1, 2026 in English and were translated on February 9, 2026." Strategy: documentary dissolution · illegal, publicly admitted. Common legal basis: French Toubon Law art. 5 · Constitution art. 2 · Civil Code art. 1171 · EU Data Act art. 23.

◆ MECHANISM II — CASCADING INCORPORATION BY REFERENCE

The signed contract refers to the Product Terms. The Product Terms refer to the DPA. The DPA refers to the SLA. The SLA refers to online documentation that can be unilaterally modified. The most binding clause is always several layers below the signed document. 🔷 Microsoft EES — "About this Document" (EN only): "Updates that Microsoft makes from time to time to Use Rights apply to Customer as set forth in Customer's agreement." Absent from the French version. Legal basis: Civil Code art. 1119 (terms only binding if brought to the other party's attention) · art. 1193 (unilateral modification prohibited).

◆ MECHANISM III — SELECTIVE NON-CANCELLABILITY

🇺🇸 AWS, s.5.4.2: "Savings Plans [...] are noncancellable [...] you will be charged for the duration [...] even if you terminate the Agreement." 🔷 Azure EES Savings Plan (FR+EN identical): "non-cancellable. The Customer must pay the amount invoiced for the selected duration, even if the Agreement is terminated." Legal basis: Civil Code art. 1171 · 1195 (unforeseeability) · Data Act art. 25 · German BGB §309 Nr.5 · UK UCTA 1977 s.3.

◆ MECHANISM IV — PRICE ASYMMETRY

🇺🇸 AWS, s.50.12.3, Bedrock throughput: "We may modify the pricing [...] at any time. [...] you will be charged for the duration [...] even if you terminate this Agreement." 🔵 GCP, s.2.6 (absent from FR version): "Google may change the Fees at any time [...] If Customer objects to any change in Fees, Customer may terminate [...] under Section 8.5" — but termination itself is subject to financial restrictions. Legal basis: Civil Code art. 1193 · Commercial Code L.442-1 · BGB §308 Nr.4 · UK CRA 2015 s.62.

MACHINA

Amine RAITI · CC BY-NC-SA 4.0

MECHANISMS V TO VII — SUSPENSION · WAIVER · RATCHET · EGRESS

◆ MECHANISM V — SUSPENSION WITHOUT NOTICE

🇺🇸 AWS, s.1.4: "we may remove or disable access to any Prohibited Content without prior notice." The threat of suspension alone is enough to discourage disputes — no need to use it, its existence is the lever. Azure MOSA s.3.c is more protective: 30-day notice for non-payment — the one point where Azure is better than AWS. Legal basis: Civil Code art. 1226 · Commercial Code L.442-1 II · UK CA 1998 s.18.

◆ MECHANISM VI — NESTED WAIVER OF LEGAL RIGHTS

🔷 Microsoft EES — clause present in EN, ABSENT from FR (Jan + June 2026): "Customer agrees to waive any and all entitlements that would otherwise be applicable under the European Electronic Communications Code (Directive 2018/1972) Article 102 paragraphs 1, 3, and 5; Article 105 paragraph 1; and Article 107 paragraphs 1 and 3." The French Ministry of Education, signatory to a €152m framework agreement, contractually waives European legal rights without that waiver being presented to it in its own language. Legal basis: Toubon Law art. 5 · Civil Code art. 1171 · Directive 2018/1972 art. 102/105/107 (public order).

◆ MECHANISM VII — THE RATCHET EFFECT

🔷 Microsoft EES, change summary, 1 June 2026: "Addition of new terms for Azure Capacity Blocks." Absent from the Jan., Feb., Apr. 2026 versions. Introduced unilaterally. Automatically applied to all existing contracts. Legal basis: Civil Code art. 1193 · art. 1119 · German BGB §305c (contractual surprise).

◆ GRADUS IV — THE PRICING LAYER: EGRESS FEES

AWS: $0.08–0.09/GB → $8,000–9,000 for 100TB — not in the ToS, in the separate price list. Google Cloud: $0.08–0.12/GB → $8,000–12,000 for 100TB — not in the ToS. Microsoft Azure: $0.07–0.09/GB → $7,000–9,000 for 100TB — partially in the ToS. For 1 petabyte: $80,000–100,000 before any technical migration cost — not contractualised. Entry into the cloud is free (ingestion = 0). Exit is charged. The machine carefully separates the contractual from the pricing layer to maximise opacity at the point of commitment. Data Act art. 25: clauses obstructing a change of provider are unenforceable.

MACHINA

Amine RAITI · CC BY-NC-SA 4.0

THE TECHNICAL LAYER AND SYSTEMIC CONCLUSION

◆ GRADUS V — THE TECHNICAL LAYER

Proprietary APIs: Lambda/Cloud Functions/Azure Functions = non-standardised interfaces, migration without rewrite is impossible. Managed services: BigQuery/Redshift/Azure Synapse = proprietary SQL dialects, non-portable formats — 10 years of queries becomes a multi-month migration project. Superlinear network effect: moving from 10 interconnected services costs far more than 10× the cost of migrating one alone. Each additional service raises the exit cost of the whole.

◆ GRADUS VI — WHAT DISTINGUISHES THIS FROM ORDINARY PRACTICE

Reason I — Scale: AWS+GCP+Azure = 65%+ of the global market. What is normal for one player among ten becomes abuse of a dominant position for a triopoly (German GWB §19a · UK CA 1998 s.18 · French L.442-1). Reason II — Information asymmetry: egress fees are not in the ToS. The waiver of Directive 2018/1972 is in English only. Constitutive of the model, not accidental. Reason III — Violations of positive law: Toubon Law · Civil Code art. 1171 · Data Act art. 23-25 · Directive 2018/1972. Illegal clauses do not constitute an aggressive business model — they constitute a non-compliant one. Reason IV — Deliberate design: drafted by top-tier legal teams, tested across jurisdictions, regularly amended. Is the Directive 2018/1972 waiver clause really only in English by accident?

◆ GRADUS VII — SYSTEMIC CONCLUSION

A single unfair clause is a matter for a civil judge. A system deliberately designed to create dependency, resting on illegal clauses, at the scale of a market with 65% concentration — that is a matter for the competition regulator. The German Bundeskartellamt (§19a GWB), the UK CMA (CA 1998 s.18) and the French DGCCRF (L.442-1) have the tools to address the systemic level. This document and the nineteen analyses underpinning it together form the two faces of the same case file: clause by clause, and coherent system. Both are necessary for regulators to act.

MACHINA

Amine RAITI · CC BY-NC-SA 4.0

ÉTUDE SYSTÉMIQUE · DOCUMENT PUBLIC · OPÉRATION DINDON · JUIN 2026

LA MACHINE DE GUERRE CONTRACTUELLE

The Contractual War Machine · Die Vertragliche Kriegsmaschine

MACHINA

🇫🇷 FRANÇAIS

Destinataire : **DGCCRF · CNIL · ARCEP**
Base légale : Loi Toubon · Code Civil · Data Act
Marché : France entière

🇬🇧 ENGLISH

Recipient: **CMA · FCA · ICO**
Legal basis: UCTA 1977 · CRA 2015 · CA 1998
Market: United Kingdom

🇩🇪 DEUTSCH

Empfänger: **Bundeskartellamt · BMWi**
Rechtsgrundlage: BGB §§ 305c/307 · GWB §19a
Markt: Deutschland

3 Couches · Layers · Ebenen	**7** Mécanismes · Mechanisms · Mechanismen	**19** Analyses · Legal analyses · Rechtsanalysen	**65%** Part de marché mondiale

Sources : CGV AWS (FR/EN/DE · avr.–mai 2026) · CGV GCP (FR/EN/DE · 2022–2026)
MOSA Azure (mars 2019) · Product Terms MOSA & EES (janv.–juin 2026)
Débat parlementaire n°5312 · Usine Digitale mars 2026 · CMA Cloud Market Study 2023

AMINE RAITI · Architecte Infrastructure & SRE

Paris le 10 Juin 2026 · Document public · *CC BY-NC-SA 4.0*

🇫🇷 VERSION FRANÇAISE · DESTINATAIRE : DGCCRF · CNIL · ARCEP

LA MACHINE DE GUERRE CONTRACTUELLE

Anatomie Complète · Version Française

Déconstruction systémique du vendor lock-in hyperscaler
fondée sur dix-neuf analyses juridiques des CGV AWS · GCP · Azure
à destination de la DGCCRF · CNIL · ARCEP · Régulateurs européens

MACHINA

BASE LÉGALE FRANÇAISE

- ◆ Loi Toubon n°94-665 art. 5
- ◆ Code Civil art. 1119 · 1171 · 1193 · 1195
- ◆ Code Commerce L.442-1
- ◆ Data Act (UE 2023/2854) art. 23 · 25
- ◆ Directive 2018/1972 art. 102 · 105 · 107
- ◆ Constitution art. 2
- ◆ Code Civil art. 1217 · 1226
- ◆ Code Conso L.212-1 · R.212-2
- ◆ Arrêt Chronopost 1996
- ◆ RGPD art. 5 · 6 · 13-14

Amine RAITI · Architecte Infrastructure & SRE

Document public · CC BY-NC-SA 4.0 · Pas de conseil juridique

PRINCIPE DE CONCEPTION · ARCHITECTURE TROIS COUCHES · CYCLE DE CAPTIVITÉ

⚠ CONSTAT FONDAMENTAL

Ces contrats ne sont pas mal rédigés. Ils sont **PARFAITEMENT RÉDIGÉS** pour un objectif précis qui n'est pas celui du client : **MAXIMISER LE COÛT DE SORTIE** tout en minimisant le coût d'entrée. Ce n'est pas un effet secondaire — c'est la fonction principale. Chaque clause analysée s'explique immédiatement dès qu'on lit le contrat avec cet objectif.

01 CONTRACTUEL

Non-résiliabilité · Clause de langue · Modification unilatérale · Renonciation aux droits légaux

02 TARIFAIRE

Egress fees 0,08–0,12 $/Go · Asymétrie entrée/sortie · Frais résiliation 12% (Azure)

03 TECHNIQUE

APIs propriétaires · Services managés · Dépendance superlinéaire au portefeuille

◆ PRINCIPE D'INTERACTION

Un client qui négocie un contrat résiliable reste bloqué par les egress fees. Un client qui contourne les egress fees reste bloqué par les engagements non résiliables. Un client qui construit une architecture ouverte reste bloqué par les Product Terms modifiés mensuellement. La sortie est systématiquement coûteuse, quelle que soit la stratégie.

CYCLE DE CAPTIVITÉ — 5 PHASES

ENTRE	→	S'INSTALLE	→	DÉPEND	→	S'ENGAGE	→	DÉCOUVRE
Crédits offerts · APIs excellentes · Coût nul		Services parfaits · Documentation complète		APIs dans le code · Équipes certifiées		-30% sur 3 ans · Non-résiliabilité		Trop tard · Trois couches actives

◆ LE CHOIX RATIONNEL DEVENU PIÈGE — STRUCTUREL

Les Reserved Instances coûtent 30–40% moins cher que le on-demand. La réduction est réelle. Elle s'obtient au prix de la non-résiliabilité. Le client échange sa liberté contre une réduction tarifaire — choix rationnel à court terme qui devient piège à moyen terme. **Code Civil art. 1171** : déséquilibre significatif dans un contrat d'adhésion = réputé non écrit. **Data Act art. 25** : tout obstacle contractuel au changement de fournisseur = inopposable.

Amine RAITI · CC BY-NC-SA 4.0

GRADUS III · MÉCANIQUE — MÉCANISMES I À IV

LANGUE · INCORPORATION EN CASCADE · NON-RÉSILIABILITÉ · ASYMÉTRIE DE PRIX

◆ MEC. I · CLAUSE DE LANGUE — 3 STRATÉGIES, 1 RÉSULTAT

CRITIQUE · TOUS HYPERSCALERS

AWS (en-tête de chaque page FR) :

> *"La traduction [...] est fournie à des fins informatives uniquement. [...] la version anglaise prévaudra." — Retard : 37 jours · Explicite · Illégal.*

GCP (Section 14.18, figée août 2023) :

> *"[...] la version anglaise prévaut." — Retard FR : 2 ans 9 mois · 13 MàJ manquantes · Dissimulée.*

Azure MOSA (6 couches, Product Terms mensuels) :

> *"Published February 1, 2026 in English, translated February 9, 2026." — Admis publiquement · Fragmentation documentaire.*

◆ VIOLATION

Loi Toubon art. 5 · Constitution art. 2 · Code Civil art. 1171 · Data Act art. 23

◆ MEC. II · INCORPORATION EN CASCADE — CLAUSE HORS DOCUMENT SIGNÉ

ÉLEVÉE

EES Microsoft · "About this Document" (version EN seulement) :

> *"Updates that Microsoft makes from time to time to Use Rights apply to Customer as set forth in Customer's agreement." — Absent de la version française. Opposable via incorporation par référence.*

◆ VIOLATION

Code Civil art. 1119 : conditions générales opposables seulement si portées à la connaissance de l'autre partie · **art. 1193** : modification unilatérale interdite

◆ MEC. III · NON-RÉSILIABILITÉ SÉLECTIVE — LE CŒUR DU VERROUILLAGE

CRITIQUE · AWS + AZURE

AWS Section 5.5 · Capacity Blocks :

> *"cannot be canceled nor modified [...] full price nonrefundable [...] last 30 minutes we may terminate your instances without notice."*

Azure EES · Plan d'économies (FR + EN identiques) :

> *"non annulable. Le Client devra s'acquitter du montant [...] même si le Contrat est résilié."*

◆ VIOLATION

Code Civil art. 1171 · 1195 · 1217 · Arrêt Chronopost 1996 · Code Conso L.212-1 · Data Act art. 25

◆ MEC. IV · ASYMÉTRIE DE PRIX — FOURNISSEUR MODIFIE, CLIENT BLOQUÉ

CRITIQUE

AWS Section 50.12.3 · Bedrock :

> *"modify the pricing [...] at any time. [...] charged for the duration [...] even if you terminate this Agreement."*

GCP Section 2.6 (absente version FR 2023) :

> *"Google may change the Fees at any time." — Clause inconnue des clients français mais opposable.*

◆ VIOLATION

Code Civil art. 1193 · Code Commerce L.442-1 · Code Conso R.212-2

Amine RAITI · CC BY-NC-SA 4.0

GRADUS III · MÉCANISMES V–VII + GRADUS IV · COUCHE TARIFAIRE

SUSPENSION · RENONCIATION DROITS EU · CLIQUET MENSUEL · EGRESS FEES

◆ MEC. V · SUSPENSION SANS PRÉAVIS

ÉLEVÉE

AWS Section 1.4 :

> *"remove or disable access [...] without prior notice."*

La menace de suspension suffit à décourager toute contestation. Azure MOSA Section 3.c prévoit 30 jours de préavis pour non-paiement — seul point où Azure est meilleur qu'AWS sur ce critère. **Code Civil art. 1226 · L.442-1 II Code Commerce.**

◆ MEC. VI · RENONCIATION DROITS UE — ABSENTE VERSION FRANÇAISE

CRITIQUE · EES · 152 M€ ÉDUCATION NATIONALE

EES MICROSOFT — VERSION ANGLAISE SEULEMENT (JAN + JUIN 2026) :

> *"Customer agrees to waive any and all entitlements [...] under the European Electronic Communications Code (Directive 2018/1972) Article 102 paras 1, 3, 5 ; Article 105 para 1 ; Article 107 paras 1 and 3."*

◆ VIOLATION · MARCHÉ PUBLIC 152 M€

Le ministère de l'Éducation nationale renonce à des protections légales européennes sans que cette renonciation lui soit présentée en français. **Loi Toubon art. 5 · Code Civil art. 1171 · Directive 2018/1972 art. 102/105/107** (ordre public).

◆ MEC. VII · EFFET DE CLIQUET — CONTRAT QUI S'ÉCRIT APRÈS SIGNATURE

CRITIQUE · AZURE

EES MICROSOFT · SYNTHÈSE MODIFICATIONS · 1ER JUIN 2026 :

> *"Ajout de nouvelles conditions pour les Blocs de Capacité Azure." — Absent des versions jan./fév./avr. 2026. Application automatique à tous contrats existants.*

◆ VIOLATION

Code Civil art. 1193 · art. 1119 · Accord-cadre Éducation nationale signé mars 2025 : clauses ajoutées après signature sans accord du ministère.

GRADUS IV · EGRESS FEES — LE MUR D'ENCEINTE HORS CGV

Fournisseur	Egress/Go	100 To	Dans CGV ?
AWS	0,08–0,09 $/Go	**8–9 k€**	**NON**
GCP	0,08–0,12 $/Go	**8–12 k€**	**NON**
Azure	0,07–0,09 $/Go	**7–9 k€**	**PARTIEL**
1 Pétaoctet	**80 000–100 000 € avant migration**	**Non contractualisé**	

L'entrée est gratuite, la sortie est tarifée. La machine sépare contractuel et tarifaire pour maximiser l'opacité à la signature. **Data Act art. 25** : obstacle au changement = inopposable.

Amine RAITI · CC BY-NC-SA 4.0

COUCHE TECHNIQUE · QUATRE RAISONS · TABLEAU COMPARATIF · CONCLUSION

◆ GRADUS V · COUCHE TECHNIQUE

APIs propriétaires : Lambda/Cloud Functions/Azure Functions = non standardisées. Migration impossible sans réécriture.

Services managés : BigQuery/Redshift/Azure Synapse = SQL propriétaire. 10 ans de requêtes = mois de migration.

Effet superlinéaire : chaque service supplémentaire augmente le coût de sortie de l'ensemble.

◆ GRADUS VI · 4 RAISONS QUI DISTINGUENT CELA

I — Échelle : 65%+ marché mondial → abus de position dominante. **L.442-1.**

II — Asymétrie d'information : egress fees hors CGV · renonciation Directive 2018/1972 en anglais uniquement · constitutive du modèle.

III — Droit positif violé : Loi Toubon · art. 1171 CC · Data Act art. 23-25 · Directive 2018/1972.

IV — Conception délibérée : équipes juridiques de premier rang · testés dans plusieurs juridictions · la clause Waiver n'est pas en anglais seulement par accident.

Mécanisme	AWS	GCP	Azure
Substitution linguistique	**Explicite · chaque page**	**Dissimulée s.14.18**	**Fragmentée · 6 couches**
Non-résiliabilité	**s.5.4.2 · 5.5 · 50.12.3**	**s.2.1 · engagements**	**Savings Plans · Blocs**
Remboursement si résiliation	**Prorata si AWS résilie ✓**	**Aucun · s.8.8**	**12% frais + 50k$/an**
Data Act Addendum	**✓ s.1.28**	**Absent**	**Absent MOSA/EES**
Renonciation droits UE	**Partiel**	**Partiel**	**Dir. 2018/1972 · EN seul.**

◆ GRADUS VII · CONCLUSION — CLAUSE VS SYSTÈME

Une clause abusive est un problème pour le juge civil. Un système conçu délibérément pour créer de la dépendance, s'appuyant sur des clauses illégales, à l'échelle d'un marché à 65% de concentration — c'est un problème pour le régulateur de concurrence. La **DGCCRF** dispose du **Code Commerce L.442-1**. Ce dossier et les dix-neuf analyses qui le fondent constituent les deux faces du même dossier : clause par clause, et système cohérent. Les deux sont nécessaires pour agir.

Summum ius, summa iniuria* — Cicéron · *Nemo supra legem est

Document public · CC BY-NC-SA 4.0 · Pas de conseil juridique · Amine RAITI

Amine RAITI · CC BY-NC-SA 4.0

🇬🇧 ENGLISH VERSION · RECIPIENT: CMA · FCA · ICO

THE CONTRACTUAL WAR MACHINE

Complete Anatomy · English Version

Systemic deconstruction of hyperscaler vendor lock-in
based on nineteen legal analyses of AWS · GCP · Azure terms
addressed to the CMA · FCA · ICO · European regulators

UK LEGAL BASIS

- ◆ UCTA 1977 s. 3 (reasonableness test)
- ◆ SGSA 1982 s. 13 (implied terms)
- ◆ Hadley v Baxendale [1854]
- ◆ Smith v Eric S Bush [1990]
- ◆ EU Data Act (2023/2854) art. 23 · 25
- ◆ CRA 2015 s. 62 (unfair terms)
- ◆ Competition Act 1998 s. 18
- ◆ Interfoto v Stiletto [1989]
- ◆ CMA Cloud Market Study 2023
- ◆ UK GDPR art. 5 · 6 · 13-14

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · *Not legal advice*

MACHINA

⚠ CORE FINDING

These contracts are not poorly drafted. They are **PERFECTLY DRAFTED** for a precise objective that is not the customer's: to **MAXIMISE EXIT COSTS** while minimising entry costs. This is not a side effect — it is the primary function. Every clause analysed becomes immediately explicable when read with this objective in mind rather than through the lens of a good-faith buyer.

01
CONTRACTUAL
Non-cancellability · Language clause · Unilateral modification · Waiver of legal rights

02
PRICING
Egress fees 0.08–0.12 $/GB · Entry/exit asymmetry · Early termination fees 12% (Azure)

03
TECHNICAL
Proprietary APIs · Managed services · Superlinear portfolio dependency

◆ INTERACTION PRINCIPLE

A customer who negotiates a cancellable contract remains locked in by egress fees. A customer who circumvents egress fees remains locked in by non-cancellable commitments. A customer who builds an open architecture remains locked in by monthly Product Terms updates. Exit is systematically costly regardless of strategy.

CAPTIVITY CYCLE — 5 PHASES

ENTERS — Free credits · Excellent APIs · Zero entry cost → **SETTLES** — Good services · Perfect docs → **DEPENDS** — APIs in codebase · Teams certified → **COMMITS** — −30% over 3 yrs · Non-cancellable → **DISCOVERS** — Too late · Three layers active

◆ THE RATIONAL CHOICE THAT BECAME A TRAP

STRUCTURAL

Reserved Instances cost 30–40% less than on-demand. The saving is real. It comes at the price of non-cancellability. The customer trades freedom for a price reduction — a rational short-term choice that becomes a medium-term trap. **UCTA 1977 s. 3**: unreasonable exclusion of liability = unenforceable. **CRA 2015 Sch. 2**: terms creating significant imbalance = unfair. **EU Data Act art. 25**: contractual obstacles to switching cloud providers = unenforceable.

Amine RAITI · CC BY-NC-SA 4.0

GRADUS III · MECHANICS — MECHANISMS I TO IV

LANGUAGE · CASCADE INCORPORATION · NON-CANCELLABILITY · PRICE ASYMMETRY

◆ MEC. I · LANGUAGE CLAUSE — 3 STRATEGIES, 1 RESULT

CRITICAL · ALL HYPERSCALERS

AWS (header of every page — FR version):

> *"The following translation is provided for informational purposes only. [...] the English version will prevail." — Gap: 37 days · Explicit · Unlawful.*

GCP (Section 14.18, frozen August 2023):

> *"[...] the English text will govern." — FR gap: 2 yrs 9 mths · 13 missed updates · Concealed.*

Azure MOSA (6-layer fragmentation, monthly Product Terms):

> *"Published February 1, 2026 in English, translated February 9, 2026." — Self-admitted · Documentary dissolution.*

◆ LEGAL BASIS

Interfoto v Stiletto [1989]: onerous terms must be drawn to attention · **CRA 2015 Sch. 2 para. 1(1)** · **EU Data Act art. 23**

◆ MEC. II · CASCADE INCORPORATION — CLAUSE OUTSIDE SIGNED DOCUMENT

HIGH

EES Microsoft · "About this Document" (EN version only):

> *"Updates that Microsoft makes from time to time to Use Rights apply to Customer as set forth in Customer's agreement." — Absent from French version. Enforceable via incorporation by reference.*

◆ LEGAL BASIS

Interfoto v Stiletto [1989] · **UCTA 1977 s. 3**: customer must be given fair notice of terms · **CRA 2015 s. 68**: terms must be transparent and prominent

◆ MEC. III · SELECTIVE NON-CANCELLABILITY — THE CORE LOCK-IN

CRITICAL · AWS + AZURE

AWS Section 5.5 · Capacity Blocks:

> *"cannot be canceled nor modified [...] full price nonrefundable [...] last 30 minutes we may terminate your instances without notice."*

Azure EES · Savings Plan (identical FR + EN):

> *"noncancellable. Customer will owe the amount charged [...] even if the Agreement is terminated."*

◆ LEGAL BASIS

UCTA 1977 s. 3: reasonableness test · **Hadley v Baxendale [1854]** · **CRA 2015 Sch. 2** · **EU Data Act art. 25**

◆ MEC. IV · PRICE ASYMMETRY — SUPPLIER MODIFIES, CUSTOMER LOCKED

CRITICAL

AWS Section 50.12.3 · Bedrock:

> *"modify the pricing [...] at any time. [...] charged for the duration [...] even if you terminate this Agreement."*

GCP Section 2.6 (absent from FR version 2023):

> *"Google may change the Fees at any time." — Unknown to French customers but enforceable.*

◆ LEGAL BASIS

UCTA 1977 s. 3 · **CRA 2015 s. 62** · **Smith v Eric S Bush [1990]** · **CA 1998 s. 18** (abuse of dominant position)

Amine RAITI · CC BY-NC-SA 4.0

SUSPENSION · RIGHTS WAIVER · RATCHET EFFECT · EGRESS FEES

◆ MEC. V · SUSPENSION WITHOUT NOTICE — HIGH

AWS Section 1.4:

> *"remove or disable access [...] without prior notice."*

The threat of suspension is sufficient to deter any challenge. Azure MOSA Section 3.c provides 30 days' notice for non-payment — the one point where Azure outperforms AWS on this criterion. **CA 1998 s. 18 · CRA 2015 Sch. 2 para. 1(b)**: right to suspend must not be exercised in a disproportionate manner.

◆ MEC. VI · WAIVER OF EU RIGHTS — ENGLISH VERSION ONLY — CRITICAL · EES MICROSOFT

EES Microsoft — English version only (Jan + Jun 2026):

> *"Customer agrees to waive any and all entitlements [...] under the European Electronic Communications Code (Directive 2018/1972) Article 102 paras 1, 3, 5 ; Article 105 para 1 ; Article 107 paras 1 and 3."*

◆ LEGAL BASIS

The French Ministry of Education waives EU legal protections without this waiver being presented to it in its own language. **Interfoto v Stiletto [1989]**: onerous clauses must be explicitly drawn to the other party's attention before or at the time of contracting.

◆ MEC. VII · RATCHET EFFECT — CONTRACT WRITTEN AFTER SIGNING — CRITICAL · AZURE

EES Microsoft · Summary of Changes · June 1, 2026:

> *"New terms added for Azure Capacity Blocks." — Absent from Jan/Feb/Apr 2026 versions. Applied automatically to all existing contracts.*

◆ LEGAL BASIS

UCTA 1977 s. 3 · Interfoto v Stiletto [1989] · CMA Cloud Market Study 2023: identified switching barriers as primary competition concern.

GRADUS IV · EGRESS FEES — THE WALL OUTSIDE THE CONTRACT

Provider	Egress/GB	100 TB	In T&Cs?
AWS	$0.08–0.09/GB	**£7–8k**	**NO**
GCP	$0.08–0.12/GB	**£7–10k**	**NO**
Azure	$0.07–0.09/GB	**£6–8k**	**PARTIAL**
1 Petabyte	**£70,000–90,000 before technical migration**	**Not contractualised**	

Entry is free, exit is priced. The machine deliberately separates contractual from pricing to maximise opacity at signature. **EU Data Act art. 25 · CA 1998 s. 18**: obstacle to switching = potentially unenforceable.

Amine RAITI · CC BY-NC-SA 4.0

TECHNICAL LAYER · FOUR REASONS · COMPARATIVE TABLE · CONCLUSION

◆ GRADUS V · TECHNICAL LAYER

PROPRIETARY APIS: Lambda/Cloud Functions/Azure Functions = non-standardised. Migration requires full rewrite.
MANAGED SERVICES: BigQuery/Redshift/Azure Synapse = proprietary SQL dialects. 10 years of queries = months of migration.
SUPERLINEAR EFFECT: each additional service increases the cost of exiting all others.

◆ GRADUS VI · 4 REASONS THIS IS DIFFERENT

I — SCALE: 65%+ global market → potential abuse of dominance. **CA 1998 s. 18**.
II — INFORMATION ASYMMETRY: egress fees outside T&Cs · Directive 2018/1972 waiver in English only · constitutive, not accidental.
III — POSITIVE LAW VIOLATED: UCTA 1977 · CRA 2015 · CA 1998 s.18 · EU Data Act art. 23–25.
IV — DELIBERATE DESIGN: drafted by top-tier legal teams · tested across jurisdictions · the waiver clause is not in English only by accident. CMA 2023 Study identified switching barriers — no enforcement followed.

MECHANISM	AWS	GCP	AZURE
Language substitution	Explicit · every page	Concealed s.14.18	Fragmented · 6 layers
Non-cancellability	s.5.4.2 · 5.5 · 50.12.3	s.2.1 · commitments	Savings Plans · Blocks
Refund on termination	Pro-rata if AWS terminates ✓	None · s.8.8	12% fee + $50k/yr cap
Data Act Addendum	✓ s.1.28	Absent	Absent MOSA/EES
EU rights waiver	Partial	Partial	Dir. 2018/1972 · EN only

◆ GRADUS VII · CONCLUSION — CLAUSE VS SYSTEM

An unfair clause is a problem for the civil courts. A system deliberately designed to create dependency, relying on unlawful clauses, at the scale of a 65% concentrated market — that is a problem for the competition regulator. The **CMA** holds the **COMPETITION ACT 1998 s. 18**. Its own 2023 Cloud Market Study identified switching barriers as a primary concern. No formal enforcement has followed. This document and the nineteen analyses underpinning it constitute the two faces of the same case: clause by clause, and as a coherent system. Both are necessary for regulators to act.

NEMO SUPRA LEGEM EST · SUMMUM IUS, SUMMA INIURIA — CICERO

Public document · CC BY-NC-SA 4.0 · Not legal advice · Amine RAITI

🇩🇪 DEUTSCHE FASSUNG · EMPFÄNGER: BUNDESKARTELLAMT · BMWI

DIE VERTRAGLICHE KRIEGSMASCHINE

Vollständige Anatomie · Deutsche Fassung

Systemische Dekonstruktion des Hyperscaler-Vendor-Lockins
auf Grundlage von neunzehn Rechtsanalysen der AGBs von AWS · GCP · Azure
gerichtet an Bundeskartellamt · BMWi · Europäische Regulierungsbehörden

MACHINA

DEUTSCHE RECHTSGRUNDLAGE

- ◆ BGB § 305c (überraschende Klauseln)
- ◆ BGB § 307 Abs. 1 S. 2 (Transparenzgebot)
- ◆ BGB § 309 Nr. 5 · 7 · 8b (Klauselverbote)
- ◆ GWB § 19 / § 19a (Marktmacht · GAIA)
- ◆ EU Data Act (2023/2854) Art. 23 · 25
- ◆ BGB § 307 Abs. 1 (unangemessene Benachteiligung)
- ◆ BGB § 308 Nr. 4 (Leistungsänderung)
- ◆ BGB § 314 (außerordentliche Kündigung)
- ◆ GG Art. 103 (rechtliches Gehör)
- ◆ DSGVO Art. 5 · 6 · 13-14

AMINE RAITI · Infrastrukturarchitekt & SRE

Öffentliches Dokument · CC BY-NC-SA 4.0 · Kein Rechtsrat

GESTALTUNGSPRINZIP · DREI-SCHICHTEN-ARCHITEKTUR · GEFANGENSCHAFTSZYKLUS

⚠ GRUNDLEGENDE FESTSTELLUNG

Diese Verträge sind nicht schlecht formuliert. Sie sind **perfekt formuliert** — für ein präzises Ziel, das nicht das des Kunden ist: die **Austrittskosten zu maximieren** und gleichzeitig die Eintrittskosten zu minimieren. Dies ist kein Nebeneffekt — es ist die Hauptfunktion. Jede analysierte Klausel erklärt sich sofort, wenn man den Vertrag mit diesem Ziel liest.

01
VERTRAGLICH
Unkündbarkeit · Sprachklausel · Einseitige Änderung · Rechtsverzicht

02
TARIFLICH
Egress-Gebühren 0,08–0,12 $/GB · Eintrags-/Austrittsasymmetrie · Kündigungsgebühr 12% (Azure)

03
TECHNISCH
Proprietäre APIs · Managed Services · Superlineare Portfolio-Abhängigkeit

◆ WECHSELWIRKUNGSPRINZIP

Ein Kunde, der einen kündbaren Vertrag aushandelt, bleibt durch Egress-Gebühren gebunden. Wer Egress-Gebühren umgeht, bleibt durch unkündbare Verpflichtungen gebunden. Wer eine offene Architektur aufbaut, bleibt durch monatlich geänderte Product Terms gebunden. Der Austritt ist unabhängig von der gewählten Strategie systematisch kostspielig.

GEFANGENSCHAFTSZYKLUS — 5 PHASEN

EINTRITT
Kostenlos · Exzellente APIs
→
EINRICHTUNG
Gute Services · Perfekte Doku
→
ABHÄNGIG
APIs im Code · Zertifizierungen
→
ENGAGEMENT
−30% über 3 J. · Unkündbar
→
ENTDECKUNG
Zu spät · Drei Schichten

◆ DIE RATIONALE ENTSCHEIDUNG ALS FALLE — STRUKTURELL

Reserved Instances kosten 30–40% weniger als On-Demand. Die Ersparnis ist real. Sie wird durch Unkündbarkeit erkauft. Der Kunde tauscht Freiheit gegen einen Preisnachlass — eine rational erscheinende kurzfristige Entscheidung, die mittelfristig zur Falle wird. **BGB § 307 Abs. I**: unangemessene Benachteiligung = unwirksam. **BGB § 309 Nr. 5**: Klauselverbote ohne Wertungsmöglichkeit. **EU Data Act Art. 25**: Klauseln, die den Anbieterwechsel behindern = nicht anwendbar.

Amine RAITI · CC BY-NC-SA 4.0

SPRACHKLAUSEL · KASKADEN-EINBEZIEHUNG · UNKÜNDBARKEIT · PREISASYMMETRIE

◆ MECH. I · SPRACHKLAUSEL — 3 STRATEGIEN, 1 ERGEBNIS

KRITISCH · ALLE HYPERSCALER

AWS (Kopfzeile jeder Seite — DE-Fassung):

> *"Die folgende Übersetzung dient ausschließlich Ihrer Information. [...] ist die englische Fassung maßgeblich." — Verzug: 37 Tage · Explizit · Rechtswidrig.*

GCP (Abschnitt 14.18, eingefroren August 2023):

> *"[...] gilt der englische Text, sofern in der Übersetzung nicht ausdrücklich etwas anderes angegeben ist." — DE-Verzug: 4 Jahre · 16 fehlende Updates · Verborgen.*

Azure MOSA (6 Schichten, monatliche Product Terms):

> *"Published February 1, 2026 in English, translated February 9, 2026." — Von Microsoft selbst dokumentiert.*

◆ RECHTSGRUNDLAGE

BGB § 305c: überraschende Klausel = nicht einbezogen · **§ 307 Abs. I S. 2**: Transparenzgebot · **EU Data Act Art. 23**

◆ MECH. II · KASKADEN-EINBEZIEHUNG — KLAUSEL AUSSERHALB UNTERZEICHNETEM DOKUMENT

HOCH

EES Microsoft · "About this Document" (nur EN-Fassung):

> *"Updates that Microsoft makes from time to time to Use Rights apply to Customer as set forth in Customer's agreement." — In der deutschen Fassung nicht vorhanden. Durchsetzbar via Einbeziehung durch Verweis.*

◆ RECHTSGRUNDLAGE

BGB § 305 Abs. 2: AGB nur wirksam, wenn ausdrücklich auf sie hingewiesen wurde · **§ 305c**: überraschende Klausel = nicht einbezogen

◆ MECH. III · SELEKTIVE UNKÜNDBARKEIT — DER KERN DES LOCKINS

KRITISCH · AWS + AZURE

AWS Abschnitt 5.5 · Capacity Blocks:

> *"cannot be canceled nor modified [...] full price nonrefundable [...] last 30 minutes we may terminate your instances without notice."*

Azure EES · Savings Plan (DE + EN identisch):

> *"noncancellable. Customer will owe the amount charged [...] even if the Agreement is terminated."*

◆ RECHTSGRUNDLAGE

BGB § 309 Nr. 5: unzulässige Pauschalschadenersatzklausel · **§ 309 Nr. 8b**: unzulässiger Ausschluss von Gewährleistungsrechten · **§ 314**: unveräußerliches Kündigungsrecht aus wichtigem Grund

◆ MECH. IV · PREISASYMMETRIE — ANBIETER ÄNDERT, KUNDE GEBUNDEN

KRITISCH

AWS Abschnitt 50.12.3 · Bedrock:

> *"modify the pricing [...] at any time. [...] charged for the duration [...] even if you terminate."*

GCP Abschnitt 2.6 (in DE-Fassung 2023 nicht vorhanden):

> *"Google may change the Fees at any time." — Deutschen Kunden unbekannt, aber anwendbar.*

◆ RECHTSGRUNDLAGE

BGB § 308 Nr. 4: Leistungsänderungsklausel unwirksam · **§ 307 Abs. I**: unangemessene Benachteiligung · **GWB § 19a**: Verhaltensanforderungen für Unternehmen mit überragender Marktbedeutung

Amine RAITI · CC BY-NC-SA 4.0

SPERRUNG · RECHTSVERZICHT · RATSCHENEFFEKT · EGRESS-GEBÜHREN

◆ MECH. V · SPERRUNG OHNE VORANKÜNDIGUNG — HOCH

AWS Abschnitt 1.4:

> *"remove or disable access [...] without prior notice."*

Die Sperrungsdrohung reicht aus, um jede Anfechtung zu unterbinden. Azure MOSA Abschnitt 3.c sieht 30 Tage Vorlaufzeit bei Nichtzahlung vor — der einzige Punkt, in dem Azure besser abschneidet als AWS. **BGB § 307 Abs. 1** · **GWB § 19 Abs. 2 Nr. 1**: Missbrauch einer marktbeherrschenden Stellung.

◆ MECH. VI · VERZICHT AUF EU-RECHTE — NUR IN ENGLISCHER FASSUNG — KRITISCH · EES MICROSOFT

EES Microsoft — nur englische Fassung (Jan + Jun 2026):

> *"Customer agrees to waive any and all entitlements [...] under the European Electronic Communications Code (Directive 2018/1972) Article 102 paras 1, 3, 5 ; Article 105 para 1 ; Article 107 paras 1 and 3."*

◆ RECHTSGRUNDLAGE

Das französische Bildungsministerium verzichtet auf EU-Rechtsschutz, ohne dass dieser Verzicht in seiner Sprache dargestellt wird. **BGB § 305c**: Klausel, die den anderen Teil überrascht = nicht einbezogen · **GG Art. 103**: rechtliches Gehör.

◆ MECH. VII · RATSCHENEFFEKT — VERTRAG NACH UNTERZEICHNUNG GESCHRIEBEN — KRITISCH · AZURE

EES Microsoft · Änderungsübersicht · 1. Juni 2026:

> *"New terms added for Azure Capacity Blocks." — Nicht in Jan/Feb/Apr 2026 enthalten. Automatische Anwendung auf alle bestehenden Verträge.*

◆ RECHTSGRUNDLAGE

BGB § 305c: überraschende Klausel · **§ 308 Nr. 4**: unzulässige einseitige Leistungsänderung · **BGH XI ZR 140/03**: Änderungsklausel bei Dauerschuldverhältnissen.

GRADUS IV · EGRESS-GEBÜHREN — DIE MAUER AUSSERHALB DES VERTRAGS

Anbieter	Egress/GB	100 TB	In AGBs?
AWS	0,08–0,09 $/GB	**7–8 T€**	**NEIN**
GCP	0,08–0,12 $/GB	**7–10 T€**	**NEIN**
Azure	0,07–0,09 $/GB	**6–8 T€**	**TEILWEISE**
1 Petabyte	**70.000–100.000 € vor technischer Migration**	**Nicht vertraglich**	

Eintritt kostenlos, Austritt kostenpflichtig. Die Maschine trennt Vertragliches und Tarifliches bewusst, um Intransparenz zu maximieren. **EU Data Act Art. 25** · **GWB § 19a Abs. 2 Nr. 4**: Wechselhindernisse = unzulässig.

Amine RAITI · CC BY-NC-SA 4.0

TECHNISCHE SCHICHT · VIER GRÜNDE · VERGLEICHSTABELLE · SCHLUSSFOLGERUNG

◆ GRADUS V · TECHNISCHE SCHICHT

Proprietäre APIs: Lambda/Cloud Functions/Azure Functions = nicht standardisiert. Migration erfordert vollständige Neuentwicklung.
Managed Services: BigQuery/Redshift/Azure Synapse = proprietäre SQL-Dialekte. 10 Jahre Abfragen = Monate Migration.
Superlinearer Effekt: jeder weitere Dienst erhöht die Gesamtausstiegskosten überproportional.

◆ GRADUS VI · 4 GRÜNDE FÜR DIE BESONDERHEIT

I — Marktmacht: 65%+ Weltmarkt → möglicher Marktmachtmissbrauch. **GWB § 19 / § 19a.**
II — Informationsasymmetrie: Egress-Gebühren außerhalb der AGBs · Direktiven-Verzicht nur auf Englisch · systemimmanent, nicht zufällig.
III — Verstoss gegen positives Recht: BGB §§ 305c/307/308/309 · GWB § 19a · EU Data Act Art. 23–25 · DSGVO.
IV — Bewusste Gestaltung: erstklassige Rechtsteams · in mehreren Jurisdiktionen erprobt · Die Waiver-Klausel steht nicht zufällig nur auf Englisch.

Mechanismus	AWS	GCP	Azure
Sprachsubstitution	**Explizit · jede Seite**	**Verborgen s.14.18**	**Fragmentiert · 6 Schichten**
Unkündbarkeit	**§ 5.4.2 · 5.5 · 50.12.3**	**§ 2.1 · Verpflichtungen**	**Savings Plans · Blöcke**
Erstattung bei Kündigung	**Prorata wenn AWS kündigt ✓**	**Keine · § 8.8**	**12% Gebühr + 50k$/J.**
Data Act Addendum	**✓ § 1.28**	**Fehlt**	**Fehlt in MOSA/EES**
EU-Rechtsverzicht	**Teilweise**	**Teilweise**	**Dir. 2018/1972 · nur EN**

◆ GRADUS VII · SCHLUSSFOLGERUNG — KLAUSEL VS SYSTEM

Eine missbräuchliche Klausel ist ein Problem für den Zivilrichter. Ein bewusst darauf ausgelegtes System, das Abhängigkeit schafft, sich auf rechtswidrige Klauseln stützt und auf einem 65%igen Marktanteil beruht — das ist ein Problem für die Wettbewerbsbehörde. Das **Bundeskartellamt** verfügt über **GWB § 19a** (Verhaltensanforderungen für Unternehmen mit überragender Marktbedeutung). Diese Studie und die neunzehn Analysen, auf die sie sich stützt, bilden die zwei Seiten desselben Dossiers: Klausel für Klausel und als kohärentes System. Beide sind erforderlich, damit die Behörden handeln.

Summum ius, summa iniuria — Cicero · Niemand steht über dem Gesetz

Öffentliches Dokument · CC BY-NC-SA 4.0 · Kein Rechtsrat · Amine RAITI

LEGAL ANALYSIS · PUBLIC DOCUMENT

AWS SERVICE TERMS ENGLISH VERSION (MAY 2026) VS BRITISH LAW

(hereinafter "Terms" · Standard Terms and Conditions within the meaning of UCTA 1977 and CRA 2015)

The English version is the sole contractually binding version.
Several clauses conflict with UCTA 1977, CRA 2015,
the Competition Act 1998 and UK GDPR.

1	11	6	0
Binding version	Clauses analysed	UK statutes invoked	Proceedings pending

Source EN: aws.amazon.com/service-terms/ · Last updated: **22 May 2026**

UK Switching Addendum: expressly incorporated under Section 1.30

The AWS Service Terms constitute standard terms and conditions
subject to UCTA 1977, CRA 2015, and the Competition Act 1998.

Amine Raiti

Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · Not legal advice
Intended for submission to the CMA · ICO · Ofcom

CENTRAL FINDINGS · CRITICAL

UNFAIR TERMS · EXCLUSION OF LIABILITY · LOCK-IN

AWS Service Terms (22 May 2026) · Standard Terms subject to UCTA 1977 · CRA 2015

⚠ FUNDAMENTAL FINDING

The AWS Service Terms are the sole contractually binding version. Unlike the French and German analyses, there is no translation gap to examine — every clause is directly enforceable in its English form. The analysis therefore focuses on the substantive lawfulness of these terms against British statute and case law. Five categories of violations emerge: (1) exclusion or restriction of liability contrary to **UCTA 1977**; (2) unfair terms contrary to **CRA 2015 Schedule 2**; (3) clauses restricting switching contrary to the **UK Cloud Switching Addendum** (s. 1.30) and **Competition Act 1998 s. 18**; (4) data processing clauses contrary to **UK GDPR**; (5) terms contrary to the **Supply of Goods and Services Act 1982**. Notably, AWS explicitly incorporates a UK Switching Addendum (s. 1.30) — a significant advantage over GCP — but this does not cure the underlying contractual defects.

◆ CLAUSE 1 · NON-CANCELLABLE RESERVED INSTANCES · Section 5.4.2

CRITICAL · UCTA 1977 · CRA 2015

🇬🇧 **AWS SERVICE TERMS EN (22 May 2026) — Section 5.4.2 :**

> *"Savings Plans, EC2 Reserved Instances and EC2 Dedicated Host Reservations are noncancellable [...] you will be charged for the duration of the term you selected, even if you terminate the Agreement. All amounts paid [...] are nonrefundable, except that if we terminate the Agreement other than for cause [...] we will refund you a pro-rata portion of any upfront fees paid."*

🏛 **BRITISH LAW — APPLICABLE PROVISIONS :**

> UCTA 1977 s. 3: *Where a party deals on the other's written standard terms, that party cannot restrict its liability for breach or render a substantially different performance, unless the clause satisfies the reasonableness test (s. 11 + Sch. 2).* CRA 2015 s. 62 + Sch. 2 para. 1(e): *A term requiring a consumer who fails to fulfil their obligation to pay a disproportionately high sum as compensation is unfair.*

◆ CONFLICT · PARTIAL — PRO-RATA EXCEPTION NOTED

The noncancellable nature of Reserved Instances engages **UCTA 1977 s. 3** — AWS renders a substantially different performance (charges persist even after termination) unless the reasonableness test is satisfied. Under the **CRA 2015 Sch. 2 para. 1(d)**, a term permitting a trader to retain sums paid where the trader terminates the contract is presumptively unfair. The pro-rata refund provision (where AWS terminates without cause) partially mitigates this — and constitutes a significant advantage over GCP. However, no equivalent protection exists where the customer terminates due to AWS's own material breach. **Interfoto Picture Library v Stiletto [1989]**: onerous terms not fairly brought to attention may be unenforceable.

◆ CLAUSE 2 · CAPACITY BLOCKS · Section 5.5

CRITICAL · NO REFUND · NO EXCEPTION

🇬🇧 **AWS SERVICE TERMS EN — Section 5.5 :**

> *"AWS Capacity Blocks cannot be canceled nor can they be modified, and the full price of a Capacity Block is nonrefundable. [...] During the last 30 minutes of a Capacity Block, we may terminate your instances without notice and prevent the launch of new instances in your reservation."*

🏛 **BRITISH LAW :**

> UCTA 1977 s. 3(2)(b): *AWS claims the right to render a performance substantially different from that reasonably expected (terminating instances in the last 30 minutes of a fully paid, nonrefundable block) — only valid if reasonable.* Law Reform (Frustrated Contracts) Act 1943: *where a contract is frustrated, sums already paid are recoverable. A "nonrefundable" clause cannot override this statutory right.*

◆ CONFLICT · MORE SERIOUS THAN 5.4.2 · NO PRO-RATA EXCEPTION WHATSOEVER

Unlike s. 5.4.2, there is no pro-rata refund even where AWS terminates. The combination of full non-refundability plus AWS's right to terminate the last 30 minutes without notice fails the **UCTA 1977 s. 11 reasonableness test** on its face. **CRA 2015 s. 62**: a term that causes significant imbalance in parties' rights to the detriment of the consumer is unfair. **Hadley v Baxendale [1854]**: AWS's unilateral termination without notice during an active, prepaid reservation period constitutes a foreseeable breach causing direct loss — no contractual term can exclude this liability where it results from negligence (**UCTA 1977 s. 2(1)**).

Amine RAITI · CC BY-NC-SA 4.0

ANALYSIS · CLAUSES 3, 4, 5 & 6 · CRITICAL / HIGH

LIABILITY · SUSPENSION · BEDROCK · JURISDICTION

◆ CLAUSE 3 · EXCLUSION OF LIABILITY · SECTION 5.8 + BETA S. 2.6

CRITICAL · UCTA 1977 S. 2

🇬🇧 **AWS SERVICE TERMS EN — SECTION 5.8 :**

> "IN CONJUNCTION WITH THE USE OF AMAZON EC2, YOU AGREE THAT YOUR AMAZON EC2 RESOURCES MAY BE TERMINATED OR REPLACED DUE TO FAILURE, RETIREMENT, OR OTHER AWS REQUIREMENTS."

🇬🇧 **SECTION 2.6 (BETA SERVICES) :**

> "AWS AND ITS AFFILIATES AND LICENSORS [...] DISCLAIM ALL WARRANTIES [...] INCLUDING ANY WARRANTIES THAT THE BETA SERVICES WILL BE UNINTERRUPTED, ERROR-FREE OR HARMFUL-COMPONENT FREE."

🏛 **BRITISH LAW :**

> UCTA 1977 s. 2(1): *A person cannot by reference to any contract term exclude or restrict his liability for death or personal injury resulting from negligence.* s. 2(2): *In the case of other loss, liability for negligence can only be excluded if reasonable.* Supply of Goods and Services Act 1982 s. 13: *implied term that services will be carried out with reasonable care and skill — non-excludable against consumers* (CRA 2015 s. 49).

◆ CONFLICT · BLANKET DISCLAIMER FAILS UCTA REASONABLENESS TEST

AWS's all-caps disclaimer of warranties and its right to terminate EC2 resources without meaningful restriction engage **UCTA 1977 s. 2(2)**. The disclaimer of all warranties for Beta Services — including those paid for — is subject to the reasonableness test. A blanket exclusion covering AWS's own negligence is unlikely to satisfy **UCTA 1977 Sch. 2** (bargaining positions, inducements, customer's knowledge). **Smith v Eric S Bush [1990]** (HL): exclusions of liability for negligence must satisfy the reasonableness requirement strictly.

◆ CLAUSE 4 · SUSPENSION WITHOUT NOTICE · SECTION 1.4

CRITICAL · UCTA 1977 · CRA 2015

🇬🇧 **AWS SERVICE TERMS EN — SECTION 1.4 :**

> *"we may remove or disable access to any Prohibited Content without prior notice in connection with illegal content, where the content may disrupt or threaten the Services or in compliance with applicable law or any governmental, regulatory or court order or request."*

🏛 **BRITISH LAW :**

> CRA 2015 s. 62 + Sch. 2 para. 1(b): *unfair if it has the effect of permitting the trader to terminate the contract on discretionary grounds while the consumer remains bound.* Competition Act 1998 s. 18: *terms giving a dominant undertaking the unilateral right to suspend service without objective justification may constitute an abuse of dominant position.*

◆ CONFLICT · UNILATERAL SUSPENSION ON AWS'S OWN ASSESSMENT

The provision grants AWS unilateral, subjective discretion to suspend services — including where content is deemed merely to "threaten" services, a standard AWS itself defines. **CRA 2015 Sch. 2 para. 1(b)**: allowing a trader to terminate at will while the consumer remains bound is a presumptively unfair term. **Competition Act 1998 s. 18**: if AWS holds a dominant position in the relevant cloud services market (a contested but plausible finding given CMA market studies), unilateral suspension without objective criteria may constitute an abuse.

Amine RAITI · CC BY-NC-SA 4.0

LIABILITY · SUSPENSION · BEDROCK · JURISDICTION

◆ CLAUSE 5 · BEDROCK LOCK-IN · SEC. 50.12.3

CRITICAL · CRA 2015

🇬🇧 **SECTION 50.12.3 :**

> *"Provisioned throughput commitments are not transferable or refundable. Therefore, you will be charged for the duration of the term you selected even if you terminate this Agreement."*

🏛 **BRITISH LAW :**

> *CRA 2015 Sch. 2 para. 1(e): requiring a consumer who fails to fulfil their obligations to pay a disproportionately high sum. AWS may modify pricing at any time for future commitments while existing commitments remain locked and non-refundable.*

◆ CONFLICT · MAXIMUM IMBALANCE

AWS reserves the unilateral right to modify pricing while the customer has no exit right. **CRA 2015 s. 62**: significant imbalance to consumer's detriment = unfair. **UK Cloud Switching Addendum (s. 1.30)**: expressly incorporated — but does not address this pricing asymmetry.

◆ CLAUSE 6 · JURISDICTION · AWS CUSTOMER AGREEMENT

HIGH · CONFLICT OF LAWS

🇬🇧 **AWS CUSTOMER AGREEMENT — GOVERNING LAW :**

> *The AWS Customer Agreement (incorporated by reference into the Service Terms) specifies that disputes involving AWS EMEA SARL are governed by Luxembourg law and courts – for UK customers post-Brexit, this raises direct conflict-of-laws issues.*

🏛 **BRITISH LAW :**

> *UK Consumer Rights Act 2015 s. 65: liability for death or personal injury from negligence cannot be excluded. Contractual choice of foreign law cannot deprive a UK consumer of protections afforded by UK mandatory law.*

◆ CONFLICT · POST-BREXIT EXPOSURE

Post-Brexit, EU Regulation Brussels I bis no longer applies in the UK. UK consumers retain protection under **Private International Law (Miscellaneous Provisions) Act 1995** and **Rome I Regulation (retained)**: mandatory UK consumer protections cannot be displaced by a choice of Luxembourg law.

Amine RAITI · CC BY-NC-SA 4.0

ANALYSIS · CLAUSES 7, 8, 9, 10 & 11 · HIGH / MEDIUM / COMPLIANT

DATA · COMPETITION · TRAINING · SWITCHING · ADDENDUM

◆ CLAUSE 7 · DATA PROCESSING — USAGE DATA · SECTION 1.20

HIGH · UK GDPR

🇬🇧 **AWS SERVICE TERMS EN — SECTION 1.20 :**

> *"We may use information about your use of and interactions with the Services to improve the Services."*

🇬🇧 **SECTION 50.3 (AI SERVICES — CODEGURU, COMPREHEND, LEX, POLLY, REKOGNITION, TEXTRACT, TRANSCRIBE, TRANSLATE) :**

> *"(a) we may use and store AI Content [...] to develop and improve the applicable Amazon AI Service and its underlying technologies; (b) we may use and store AI Content that is not personal data to develop and improve AWS and affiliated machine learning and AI technologies."*

🏛 **BRITISH LAW :**

> *UK GDPR Art. 5(1)(b) (purpose limitation): data may only be collected for specified, explicit and legitimate purposes. "Improving services" is insufficiently specific. UK GDPR Art. 6: legitimate interest under Art. 6(1)(f) requires a documented balancing test — absent here. ICO Guidance on AI and data protection: training AI models on customer data requires explicit consent or a documented legitimate interest, not a broad contractual consent buried in standard terms.*

◆ CONFLICT · UK GDPR PURPOSE LIMITATION · ICO ENFORCEMENT RISK

The combination of s. 1.20 (general usage data) and s. 50.3 (AI Content for model training) constitutes a broad, open-ended consent to data processing for purposes that are neither specific nor explicitly stated. **UK GDPR ART. 5(1)(B)**: purpose limitation violated. **UK GDPR ART. 13**: information provided at the time of data collection must include the specific purposes — "improve the Services" does not satisfy this. The ICO has issued guidance specifically addressing AI training data and consent requirements.

◆ CLAUSE 8 · MARKET DOMINANCE · S. 50.5 + 50.6

HIGH · CA 1998 S. 18

🇬🇧 **SECTIONS 50.5 + 50.6 :**

> *"You will not use, or allow a third party to use, AI Services [...] to develop or improve a similar or competing product or service."* [s. 50.6]
> *"You may not use AI Services to generate Content for the express purpose of training an AI model."* [s. 50.5]

🏛 **BRITISH LAW :**

> *Competition Act 1998 s. 18: abuse of dominant position — restricting competition by preventing customers from using data or outputs to develop competing services. CMA has specifically investigated cloud market competition (2023 cloud market study).*

◆ POTENTIAL ABUSE · CMA SCRUTINY

Clauses preventing use of outputs for competitive development are directly within the scope of the CMA's cloud market investigation. **CA 1998 S. 18**: if AWS holds a dominant position, these restrictions on downstream competition may constitute an abuse. **TRADE MARKS ACT 1994 S. 10**: benchmarking restrictions (s. 1.8) raise additional IP law concerns.

◆ CLAUSE 9 · TRAINING FEES · SEC. 67.2.1

HIGH · CRA 2015

🇬🇧 **SECTION 67.2.1 :**

> *"Prepaid funds are nonrefundable and expire at the end of the period specified in the applicable Order."*

🏛 **BRITISH LAW :**

> *CRA 2015 s. 62 + Sch. 2 para. 1(d): a term permitting the trader to retain sums paid where the trader does not provide the service — presumptively unfair. Consumer Contracts Regulations 2013: right of withdrawal of 14 days for services not yet commenced.*

◆ CONFLICT · CRA 2015 SCH. 2

Prepaid training funds that expire with no refund even if AWS modifies or withdraws training offerings fails **CRA 2015 SCH. 2 PARA. 1(D)**. AWS reserves the right to modify training content — removing courses already paid for. **CONSUMER CONTRACTS REGS 2013 REG. 29** may apply.

Amine RAITI · CC BY-NC-SA 4.0

DATA · COMPETITION · TRAINING · SWITCHING · ADDENDUM

◆ CLAUSE 10 · UK SWITCHING ADDENDUM · SECTION 1.30

COMPLIANT · EXPRESSLY INCORPORATED

🇬🇧 **AWS SERVICE TERMS EN — SECTION 1.30 :**

> *"These Service Terms incorporate the AWS UK Switching Addendum for customers that have an eligible account, as defined in the AWS UK Switching Addendum."*

🏛 **BRITISH LAW — CONTEXT :**

> *The AWS UK Switching Addendum is a direct response to the CMA's cloud market investigation (2023). It commits AWS to providing customers with reasonable assistance to switch or multi-cloud, reducing egress fees, and providing data portability. The CMA's market study identified cloud switching barriers as a primary competition concern.*

◆ SIGNIFICANT ADVANTAGE OVER GCP · BUT ADDENDUM SCOPE IS LIMITED

This is a material positive distinction: AWS explicitly incorporates a UK switching commitment; GCP does not have an equivalent UK-specific addendum. However, the Addendum does not cure the substantive defects in ss. 5.4.2, 5.5 and 50.12.3 — noncancellable commitments continue to operate as switching barriers regardless of the Addendum's portability provisions. **Competition Act 1998 s. 18**: the CMA may still scrutinise these residual switching costs as a systemic barrier.

◆ CLAUSE 11 · SPOT INSTANCES WITHOUT NOTICE · SECTION 5.3

HIGH · UCTA 1977

🇬🇧 **AWS SERVICE TERMS EN — SECTION 5.3 :**

> *"We may terminate, stop, or hibernate Spot Instances at any time and without any notice to you if we determine that the current Spot Price [...] equals or exceeds the maximum price you specified."*

🏛 **BRITISH LAW :**

> *UCTA 1977 s. 3: AWS cannot by standard terms entitle itself to render a substantially different performance from that reasonably expected, unless the clause satisfies the reasonableness test. The market-pricing logic of Spot Instances provides some objective justification — but the extension to Spot Blocks (fixed-duration) is harder to defend.*

◆ PARTIAL CONFLICT · SPOT BLOCKS MORE PROBLEMATIC THAN SPOT INSTANCES

For true Spot Instances (variable pricing), the no-notice termination reflects the agreed market-pricing model — likely to satisfy the **UCTA 1977 s. 11 reasonableness test** in a B2B context. However, the same clause applies to **Spot Blocks** (fixed-duration purchases) — where the economic justification is weaker and the clause is more likely to fail reasonableness. **Interfoto Picture Library v Stiletto [1989]**: onerous terms in standard contracts require specific, prominent notice to be binding.

Amine RAITI · CC BY-NC-SA 4.0

COMPARATIVE TABLE · CONCLUSION · AWS vs GCP vs BRITISH LAW

Clause / Section Severity	AWS Terms (May 2026)	British Law Position	Primary Authority
s. 5.5 · Capacity Blocks — non-refundable, terminate without notice CRITICAL	**No refund · No exception**	Unlawful — fails reasonableness	UCTA 1977 s. 3 · CRA 2015 s. 62 · Hadley v Baxendale
s. 5.4.2 · Reserved Instances — noncancellable CRITICAL	Noncancellable · pro-rata if AWS terminates	Partial conflict — pro-rata mitigates	UCTA 1977 s. 3 · CRA 2015 Sch. 2 · Interfoto
s. 5.8 + s. 2.6 · Blanket liability disclaimers CRITICAL	All-caps disclaimer of all warranties	Unlawful for negligence without reasonableness	UCTA 1977 s. 2 · SGSA 1982 s. 13 · Smith v Bush
s. 1.4 · Suspension without notice CRITICAL	At AWS's sole discretion	Presumptively unfair — no objective criteria	CRA 2015 Sch. 2 para. 1(b) · CA 1998 s. 18
s. 50.12.3 · Bedrock — lock-in with unilateral pricing CRITICAL	Non-refundable · AWS modifies price at will	Unfair — maximum imbalance	CRA 2015 s. 62 · CA 1998 s. 18
s. 1.20 + s. 50.3 · Data for AI training — vague purpose HIGH	Broad consent embedded in Terms	Conflict with purpose limitation	UK GDPR Art. 5(1)(b) · Art. 6 · ICO Guidance
s. 50.5–50.6 · No competing product from AI outputs HIGH	Broad restriction on downstream use	Potential abuse of dominance	Competition Act 1998 s. 18 · CMA Cloud Study 2023
s. 67.2.1 · Prepaid training — nonrefundable HIGH	No refund if AWS modifies offering	Presumptively unfair	CRA 2015 Sch. 2 para. 1(d) · CCR 2013
s. 5.3 · Spot Blocks — terminate without notice HIGH	No notice · fixed-duration blocks	Partial — Spot Instances defensible, Spot Blocks not	UCTA 1977 s. 11 · Interfoto [1989]
s. 5.4.2 · Pro-rata refund if AWS terminates without cause MITIGATING	**Present ✓**	Compliant · advantage over GCP	UCTA 1977 s. 11 · ✅ Mitigating
s. 1.30 · UK Switching Addendum	**Expressly incorporated ✓**	Compliant · significant advantage	CA 1998 s. 18 · CMA

◆ **LEGAL CONCLUSION · AWS vs GCP · BRITISH LAW**

The AWS Service Terms, as the sole binding version, contain **two structural levels of conflict** with British law. The first is **statutory**: blanket liability disclaimers, noncancellable commitments and unilateral suspension clauses fail the reasonableness test under **UCTA 1977 s. 3 and s. 11**, and constitute presumptively unfair terms under **CRA 2015 s. 62 and Sch. 2**. The Capacity Blocks clause (s. 5.5) is the most problematic — no refund, no exception, plus unilateral termination without notice in the last 30 minutes. The second is **regulatory**: restrictions on downstream competitive use (ss. 50.5–50.6) and embedded AI training consents (ss. 1.20, 50.3) engage **Competition Act 1998 s. 18** and **UK GDPR** respectively — areas of active CMA and ICO focus. AWS is materially better than GCP in two respects: (1) the pro-rata refund mechanism where AWS terminates without cause; (2) the explicit incorporation of the **UK Switching Addendum** (s. 1.30) — a direct CMA-responsive commitment absent from GCP's terms. These are significant differentiators but do not cure the underlying substantive defects.

Source EN: aws.amazon.com/service-terms/ · 22 May 2026 · UK Switching Addendum: incorporated under s. 1.30
Statutes: UCTA 1977 · CRA 2015 · CA 1998 · SGSA 1982 · UK GDPR · Consumer Contracts Regulations 2013
Cases: Hadley v Baxendale [1854] · Interfoto v Stiletto [1989] · Smith v Eric S Bush [1990]

Nemo supra legem est

Public document · CC BY-NC-SA 4.0 · Not legal advice · Amine RAITI

LEGAL ANALYSIS · PUBLIC DOCUMENT

GOOGLE CLOUD TERMS OF SERVICE vs BRITISH LAW

9 clauses potentially void or unenforceable under current British law

4	9	0
Critical clauses	Acts violated	Proceedings active

Source: Google Cloud Terms of Service · Last modified June 1, 2026
cloud.google.com/terms · Publicly verifiable sources

Amine RAITI
Infrastructure Architect & SRE · SIPS/BCE

Public document · CC BY-NC-SA 4.0 · Free to share
This document does not constitute legal advice. It is intended for submission to competent authorities.

ANALYSIS · CLAUSES 1, 2 & 3 · CRITICAL

JURISDICTION · FEE CHANGES · NO REFUNDS

◆ CLAUSE 1 · CALIFORNIA JURISDICTION · SECTION 14.12(C)

CRITICAL

"ALL CLAIMS ARISING OUT OF OR RELATING TO THIS AGREEMENT WILL BE GOVERNED BY CALIFORNIA LAW AND WILL BE LITIGATED EXCLUSIVELY IN THE FEDERAL OR STATE COURTS OF SANTA CLARA COUNTY, CALIFORNIA, USA."

Unfair Contract Terms Act 1977 · Section 3

A business cannot impose terms on another party which are substantially unfair in a standard form contract. Forcing a British SME to litigate in California is economically prohibitive and constitutes precisely such unfairness.

Civil Jurisdiction and Judgments Act 1982 · Private International Law Act 2020

Post-Brexit, British courts maintain the principle that a jurisdiction clause in an adhesion contract cannot deprive a British business of its rights before UK courts.

◆ CLAUSE 2 · UNILATERAL FEE REVISION · SECTION 2.6

CRITICAL

"Google may change the Fees at any time unless otherwise expressly agreed in an addendum or Order Form."

Unfair Contract Terms Act 1977 · Section 3(2)(b)

A business cannot reserve to itself the right to substantially alter its contractual obligations in a standard form contract without the other party's agreement.

Consumer Rights Act 2015 · Section 62

A term that creates a significant imbalance in the parties' rights and obligations to the detriment of the consumer or micro-business is unfair and not binding.

Supply of Goods and Services Act 1982 · Section 15

Where no price is determined by mutual agreement, the price must be a reasonable one. A unilateral right to revise fees "at any time" undermines this principle.

◆ CLAUSE 3 · NO REFUNDS · SECTION 8.8

CRITICAL

"Unless expressly stated otherwise in this Agreement or required by law, termination or non-renewal under any section of this Agreement will not oblige Google to refund any Fees."

Unfair Contract Terms Act 1977 · Section 3

A business cannot exclude liability for its own non-performance or reserve the right to retain sums paid without providing the agreed service.

Consumer Rights Act 2015 · Schedule 2 · Grey List

Expressly lists as potentially unfair: terms enabling the seller to retain sums paid without delivering the corresponding service.

Law Reform (Frustrated Contracts) Act 1943

Where a contract becomes impossible to perform, sums already paid must be returned. A blanket "no refunds" clause cannot override this statutory right.

Amine RAITI · CC BY-NC-SA 4.0

ANALYSIS · CLAUSES 4, 5, 6 & 7 · CRITICAL / HIGH

PAYMENT · SUSPENSION · WARRANTIES · LIABILITY

◆ CLAUSE 4 · NON-CANCELLABLE PAYMENT OBLIGATION · SECTION 2.1 — CRITICAL

> *"Unless required by law, Customer's obligation to pay all Fees is non-cancellable."*

COMPETITION ACT 1998 · SECTION 18
Prohibits the abuse of a dominant market position. Imposing non-cancellable payment obligations on a captive market constitutes such abuse.

CONSUMER RIGHTS ACT 2015 · SECTION 62
A non-negotiated term creating a significant imbalance to the customer's detriment is unfair and not binding. An absolute, non-cancellable payment obligation — with no exit — creates precisely this imbalance.

COMMON LAW · BCCI V ALI [2001]
An absolute contractual obligation that disregards changed circumstances can be contrary to good faith principles under English Common Law.

◆ CLAUSE 5 · IMMEDIATE SUSPENSION WITHOUT RECOURSE · SECTION 4.2 — HIGH

> *"Google may immediately Suspend all or part of Customer's use of the Services if Google reasonably believes Suspension is needed to protect the Services."*

COMMON LAW · NATURAL JUSTICE · AUDI ALTERAM PARTEM
A party must be heard before a decision affecting its rights is taken. Immediate suspension on Google's unilateral assessment, without prior notice or contradictory process, violates this foundational principle.

UNFAIR CONTRACT TERMS ACT 1977 · SECTION 3(2)
A business cannot reserve the right to perform the contract in a substantially different manner from what was reasonably expected. The CMA Cloud Market Investigation 2026 specifically identified unilateral suspension practices as a concern.

◆ CLAUSE 6 · DISCLAIMER OF WARRANTIES · SECTION 11 — HIGH

> *"Google does not make and expressly disclaims to the fullest extent permitted by applicable law any warranties of any kind, including warranties of merchantability, fitness for a particular use, or error-free use."*

SUPPLY OF GOODS AND SERVICES ACT 1982 · SECTION 13
Implies a warranty that services will be carried out with reasonable care and skill — this implied term cannot be excluded in a B2B contract under UCTA 1977 · Section 2(2) unless the exclusion satisfies the reasonableness test, which a blanket disclaimer does not.

◆ CLAUSE 7 · LIABILITY CAP 12 MONTHS · SECTION 12.2 — HIGH

> *"Each party's total aggregate Liability is limited to the Fees Customer paid during the 12 month period before the event giving rise to Liability."*

UCTA 1977 · SECTION 2(2) · SECTION 11 · SCHEDULE 2
A liability limitation clause must satisfy the reasonableness test. Capping liability at 12 months of fees when a data loss could cause millions in damages fails this test.

HADLEY V BAXENDALE [1854]
Foreseeable damages must be compensable. A clause systematically excluding them contradicts this foundational Common Law principle.

Amine RAITI · CC BY-NC-SA 4.0

LANGUAGE · BRAND · FULL SYNTHESIS

◆ CLAUSE 8 · ENGLISH TEXT PREVAILS · SECTION 14.18 — MEDIUM

> *"If this Agreement is translated into any language other than English, and there is a discrepancy between the English text and the translated text, the English text will govern."*

UCTA 1977 · Interfoto Picture Library v Stiletto [1989]
Unusual or onerous clauses must be brought expressly to the attention of the contracting party. A clause granting primacy to one language version without distinct warning may be contestable on this ground. Post-Brexit UK has no Toubon-equivalent but Common Law fairness principles apply.

◆ CLAUSE 9 · USE OF CUSTOMER BRAND · SECTION 9 — MEDIUM

> *"Google may use Customer's name and Brand Features in online or offline promotional materials of the Services."*

Trade Marks Act 1994 · Section 10
Protects registered trade marks against unauthorised use. Consent buried in standard adhesion terms without separate negotiation may be insufficient.
UK GDPR · Article 7 · Data Protection Act 2018
Requires free, specific, informed and unambiguous consent. A general click-through acceptance does not satisfy this standard for commercial use of identity.

Clause	Section	British Law Violated	Severity
California jurisdiction	14.12(c)	Civil Jurisdiction Act · UCTA S.3	CRITICAL
Unilateral fee revision	2.6	UCTA S.3(2)(b) · Consumer Rights Act S.62	CRITICAL
No refunds	8.8	UCTA S.3 · Consumer Rights Act Sch.2	CRITICAL
Non-cancellable payment	2.1	Competition Act S.18 · Consumer Rights S.62	CRITICAL
Immediate suspension	4.2	Common Law Natural Justice · UCTA S.3(2)	HIGH
Disclaimer of warranties	11	Supply of Goods Act S.13 · UCTA S.2(2)	HIGH
12-month liability cap	12.2	UCTA S.2(2) · S.11 · Hadley v Baxendale	HIGH
English text prevails	14.18	UCTA · Interfoto v Stiletto 1989	MEDIUM
Use of customer brand	9	Trade Marks Act S.10 · UK GDPR Art.7	MEDIUM

◆ KEY DIFFERENCE FROM FRENCH ANALYSIS

The English version of the ToS reveals two additional critical clauses not present as explicitly in the French version — Section 2.6 (fee revision "at any time") and Section 8.8 ("No Refunds"). The English version is globally more explicit in its exclusions, making it more legally vulnerable precisely because its intent is unambiguous. The CMA's own March 2026 Cloud Market Investigation named switching barriers and fee practices as significant concerns — without triggering formal enforcement. The question is why.

Source: cloud.google.com/terms · Last modified June 1, 2026 · Public document · This document does not constitute legal advice.

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART IV

THE PHYSICAL SUBSTRATE: IRON

12 pieces

HUMAN

STRUCTURAL STUDY · CORPUS CONCLUSION · JUNE 2026

MASTERY OF THE IRON

The Final Lock and the Path to Reconquest

Sequel and conclusion to Anatomy of the Loss

RATIO

◆ FOREWORD — A REQUEST FOR READING

Like the previous document in this corpus, this study asks for a rational rather than an emotional reading. It makes no accusation of organised conspiracy. It describes a market mechanism — the concentration of a scarce skill through a lack of outside investment — and draws an operational conclusion from it: this lock, being the product not of a deliberate plan but of an investment imbalance, can be lifted through a training effort at a scale no private company, however powerful a trio, can match against the sheer number of people available worldwide.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

Operation Dindon

TABLE OF CONTENTS

- ◆ Section 1 · Mastery of the Iron as the Final Lock
- ◆ Section 2 · Concentration Through External Underinvestment
- ◆ Section 3 · What the Trio Actually Does With This Knowledge
- ◆ Section 4 · The Advantage of Numbers
- ◆ Section 5 · Conclusion — The Path to Reconquest

◆ THIS DOCUMENT'S PLACE IN THE CORPUS

This study concludes the trajectory opened by "Anatomy of the Loss — Diagnosis and Reformation." Prior reading of that document is assumed for the full coherence of the reasoning, in particular the six-layer lock-in model and the 2005 pedagogical foundation, here extended down to the hardware layer itself.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · THE FINAL LOCK

MASTERY OF THE IRON AS THE FINAL LOCK

The six-layer lock-in model already established in this corpus places the cognitive layer as the deepest. This study proposes to go one level deeper still: beneath the cognitive layer lies the hardware layer itself — the real-world knowledge of building and operating a large-scale datacentre, a distributed object storage system comparable to S3, and a global network backbone.

This knowledge is not an abstract academic skill. It is physical engineering knowledge: how to size electrical power for a multi-megawatt site, how to design a fault-tolerant storage network across hundreds of thousands of disks, how to operate a continent-scale fibre backbone. Today, this knowledge exists almost exclusively among a small number of players — the trio at the forefront, alongside a handful of legacy telecom operators and a few specialised hosting providers.

> **◆ WHY THIS LAYER IS THE MOST POWERFUL**
>
> All previous lock-in layers — contractual, pricing, technical, cognitive — can in theory be bypassed as long as the physical competence to build an alternative exists somewhere. But if the very competence to build a datacentre, a distributed storage system, or a network backbone becomes scarce enough that it only exists among the dominant players, then none of the previous layers even need to be actively maintained: the absence of outside competence alone is enough to guarantee the dominant position.

This observation is what motivates this study. If the trio's current dominance ultimately rests on the scarcity of hardware mastery rather than on contractual clauses or reversible technical choices, then reconquering that mastery is the most fundamental lever of all those identified in this corpus.

> **◆ NASSIHA — WHAT THIS OBSERVATION DOES NOT PREJUDGE**
>
> Noting that this competence is scarce does not prejudge the reason for that scarcity. That is precisely the question the next section examines, with the same demand for rigour as the rest of this corpus.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE REAL MECHANISM

CONCENTRATION THROUGH EXTERNAL UNDERINVESTMENT

The scarcity of hardware mastery requires no deliberate suppression to take hold. The mechanism is simpler, more systemic, and just as effective: if nearly all global investment in training, research and physical storage and compute infrastructure concentrates among three players, then competence concentrates mechanically among those same three players — with no need for any plan to suppress knowledge to produce that outcome.

This mechanism is exactly the one described in Section 5 of the previous document in this corpus, here applied to the scale of the global labour market rather than a single technical team. Universities, vocational training bodies and employers follow market demand. The market itself followed the path of least resistance offered by cloud abstraction. At the end of this conspiracy-free chain, outside investment in hardware competence dried up, and the competence itself became scarce through lack of practice rather than deliberate confiscation.

◆ WHY THIS NUANCE CHANGES EVERYTHING

If the scarcity resulted from a deliberate plan to suppress knowledge, reconquest would require outmanoeuvring an active strategy — an adversary that reacts and adapts. If the scarcity results from external underinvestment, reconquest requires only one thing: reinvesting. No adversarial strategy to outmanoeuvre, only an imbalance to correct. This is a more modest thesis in its accusation, and a more powerful one in its practical consequence.

◆ NASSIHA — DISTINGUISHING MECHANISM FROM ACCUSATION

This study does not claim the trio has a plan to suppress hardware knowledge across the global market. Such a claim would require specific evidence this document does not possess. It claims that the observable net effect — the concentration of competence — results from an investment imbalance between three private players and the rest of the world, an imbalance that took hold without requiring any centralised intent to explain it.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · FACTUAL HONESTY

WHAT THE TRIO ACTUALLY DOES WITH THIS KNOWLEDGE

In the interest of accuracy, it must be said that the trio does not systematically hide its hardware knowledge. Each publishes, to varying degrees, research on its infrastructure: internal chip design, datacentre cooling, distributed storage architecture. This communication serves their technical reputation and recruitment, without thereby transferring the actual operational competence required to build a complete equivalent.

◆ THE DISTINCTION THAT MATTERS

Publishing a research paper on a hardware innovation is not the same as transferring the end-to-end competence needed to build, operate and evolve a comparable system at scale. The trio keeps the essential in-house: decades of cumulative operational experience, operating processes, relationships with component manufacturers, and a critical mass of engineers trained on real production systems. It is this cumulative competence, not the absence of any publication, that constitutes the real lock.

This nuance matters because it avoids an easily refutable accusation. Claiming the trio "hides everything" would be inaccurate and would weaken the entire argument against the first Google search showing otherwise. Claiming the trio holds a rare cumulative competence, built over decades of concentrated investment, is by contrast accurate and verifiable.

◆ NASSIHA — THE TRIO IS NOT THE ENEMY OF TRANSPARENCY

Acknowledging that the trio contributes to public research on hardware infrastructure does not weaken this study's thesis. It makes it more precise: the lock is not a guarded secret, it is a cumulative competence that the scarcity of outside investment has prevented anyone else from matching.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE BALANCE OF POWER

THE ADVANTAGE OF NUMBERS

The trio consists of three private companies. The world has several million software and electronics engineers, tens of thousands of universities and engineering schools, and hundreds of countries with sovereign interests in diversifying their technological dependence. The structural balance of power does not favour the trio once competence ceases to be scarce.

This is a simple economic principle: an artificial scarcity does not need to be fought head-on, it collapses on its own once the supply of competence increases sufficiently. If mastery of the iron becomes widely available again — through the training described in the previous document of this corpus, at the scale of thousands rather than a few hundred specialised engineers — the scarcity underpinning the trio's dominant position erodes mechanically, with no frontal battle needed against an adversarial strategy that, by construction, does not exist as a centralised plan.

◆ WHAT NUMBERS CONCRETELY ENABLE

The trio today holds a handful of dominant platforms. But nothing in the nature of this competence requires it to remain held by only three players. With a few hundred engineers expert in distributed storage, backbone networking and datacentre operations, spread across multiple continents and legal frameworks, it becomes possible to bring about not a single competitor, but several independent alternatives simultaneously — precisely because the world is larger than three companies.

This reasoning connects directly to the earlier documents in this corpus: the R&D subsidy proposals for local hosting providers, the Open Hardware programme involving the military and universities, and the pedagogical foundation for hardware training. Each of these levers, taken alone, seems modest against the trio. Taken together, at the scale of competence available worldwide, they constitute a reconquest force that three private companies cannot structurally match in terms of total human capacity.

◆ NASSIHA — A NOTE OF REALISM

It would be dishonest to claim this reconquest is fast or costless. Building a credible equivalent of a global object storage system or an intercontinental network backbone requires capital, time, and overcoming network effects already established over fifteen years. The advantage of numbers is real in terms of available human competence; it does not eliminate the capital and time challenge of the physical construction itself.

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · CORPUS CONCLUSION

THE PATH TO RECONQUEST

This document concludes a line of reasoning that began with the analysis of the trio's contractual workarounds, continued with the diagnosis of behavioural, economic and cognitive mechanisms that diluted historical Ops competence, and arrives here at the deepest layer: the scarcity of mastery over the iron itself.

The practical conclusion of this trajectory is not a call to confrontation against an organised strategic adversary. It is a call to investment. The scarcity underpinning the trio's position is not an actively defended wall — it is a void that no one else has sufficiently filled. Filling that void does not require defeating a strategy, only mobilising a fraction of the world's human competence that remains, today, largely untapped on this specific terrain.

> **◆ THE LINK TO THE REST OF THE CORPUS**
>
> This document proposes nothing new that is not already sketched elsewhere in this corpus: R&D subsidies for local hosting providers, the Open Hardware programme with the military and universities, the pedagogical foundation for hardware training described in the previous document. What this conclusion adds is the framework that ties these proposals together: they are not isolated measures, they are the components of a single strategy to reconquer the most fundamental layer of the entire lock-in system.

The trio does not need to be defeated by a single player that would replace it identically. The world has the advantage of numbers, and that advantage allows for not one, but several independent alternatives, built simultaneously by engineers trained on the foundation described in this corpus, across different legal frameworks, on different continents. It is this diversity of players, more than the size of a single competitor, that constitutes the real safeguard against a return to the same lock-in under a different name.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE DIGITAL IRON

Who makes the chips on which the world's infrastructure runs

RATIO

◆ CONTEXT OF THE STUDY

This study extends "The Mastery of Iron" into its digital dimension. Physical iron — steel, cutlery, metallurgy — was treated as a condition of industrial power. Digital iron — semiconductors — is the condition of computational power. No chip, no server. No server, no infrastructure. No infrastructure, no SRE. The dependency chain runs up to a handful of actors who concentrate the manufacture of the components on which the entire global digital economy rests — and Europe is not among them.

Amine RAITI · Infrastructure Architect & SRE · Former instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE OBSERVATION

90% OF ADVANCED CHIPS MADE IN ONE PLACE

TSMC — Taiwan Semiconductor Manufacturing Company — manufactures approximately 90% of the world's most advanced logic chips (sub-7-nanometre nodes). The processors in AWS, GCP and Azure datacentres, the chips in smartphones, the NVIDIA GPUs powering generative AI, the Apple M-series processors — all come out of TSMC's factories in Taiwan. No other actor is capable of producing at this level of precision and at this scale.

This concentration is not new. It was built over forty years of massive investment, accumulated know-how passed down across generations, and a specialisation that has made competition structurally impossible for anyone who did not begin investing in the 1980s. Europe watched this concentration form without ever deciding to oppose it.

◆ GEOPOLITICAL RISK CONCENTRATED ON 35,000 KM²

Taiwan is an island of 35,000 km² located 180 km from the Chinese mainland coast. The People's Republic of China asserts sovereignty over it. A military conflict, a naval blockade, or even a major logistical disruption in the Taiwan Strait would interrupt the production of 90% of the world's advanced chips. The global digital economy — infrastructure, cloud, AI, telecoms — would face a component supply rupture within a matter of months. No buffer stock exists at this scale. No alternative production capacity is immediately available.

◆ THE CHAIN THAT RUNS DOWN TO THE SRE

The SRE administers servers. Servers contain processors. Processors are made in Taiwan. A conflict in the Taiwan Strait does not translate into an abstract geopolitical event — it translates into the impossibility of ordering new servers in the 12 to 18 months that follow, the explosion of existing component prices, and the progressive degradation of infrastructure as hardware ages with no possibility of replacement. Digital sovereignty starts with the chips. Not with the cloud contracts.

◆ NASSIHA — WHAT THIS STUDY DOES NOT CLAIM

This study does not claim that a conflict in the Taiwan Strait is imminent or likely. It claims that the concentration of advanced chip production on a single territory constitutes a documented structural risk, independently of any specific crisis scenario. A risk does not need to materialise to justify a structural response.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

DESIGN WITHOUT MANUFACTURE — THE FABLESS MODEL

The semiconductor value chain has restructured around a separation between design and manufacture. NVIDIA designs the H100 GPUs that power AI — it does not make them. Apple designs the M-series chips — it does not make them. AMD designs its EPYC datacentre processors — it does not make them. These actors are "fabless" companies: without a factory. Their designs are sent to TSMC, which etches them into silicon.

This model has produced remarkable efficiency. It allowed design companies to reach levels of sophistication they could never have funded had they needed to maintain their own production tooling. But it also created a structural vulnerability: dependence on a single manufacturer.

◆ THE DESIGN TOOLS: EQUALLY CONCENTRATED

Manufacturing is not the only bottleneck. Chip design software (EDA — Electronic Design Automation) is dominated by three American actors: Synopsys, Cadence, and Mentor Graphics (acquired by Siemens). The extreme ultraviolet (EUV) lithography machines essential for the finest nodes are produced by a single global actor: ASML, a Dutch company. Without ASML machines, TSMC cannot produce its most advanced chips. The dependency chain has multiple links, each as narrow as the last.

◆ EUROPE IN THIS CHAIN: ONE LINK

Europe is present in this chain at one significant point: ASML. This presence is real and strategic — without ASML, global advanced chip production stops. But it is also fragile: ASML itself depends on American suppliers for its lasers and subcomponents whose supply chain is not entirely European. Europe is one link in a chain it does not control end to end. STMicroelectronics, the only European chip manufacturer of scale, operates on less advanced nodes — useful for automotive and industrial applications, but not for AI datacentres.

◆ NASSIHA — ASML IS NOT AN ANSWER TO EUROPEAN DEPENDENCE

ASML is a valuable European asset. Its presence does not mean that Europe controls the semiconductor chain. Making the machines that make the chips is a position of strength — but this position depends entirely on the demand of Asian and American manufacturers. If TSMC slowed its orders, ASML would slow with it. Mastering a production tool is not the same as mastering production.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

WHAT CUTLERY TEACHES SEMICONDUCTORS

"The Mastery of Iron" documented how France allowed mastery of speciality steels to migrate to the United States and Asia — the S35VN and S45VN that make a Spyderco knife worth €450, the Hitachi steels that underpin Japanese premium cutlery. Thiers, once the world capital of cutlery, lost its pre-eminence not because it was technically defeated, but because it did not invest in metallurgical research when others were doing so.

The parallel with semiconductors is structurally identical. Europe was not excluded from advanced chip manufacturing by a hostile decision by outside actors. It excluded itself through the absence of a coherent industrial policy over forty years, while Taiwan, South Korea and the United States were investing massively.

◆ THE COMMON MECHANISM: THE ABANDONMENT OF FOUNDATIONAL INVESTMENT

In both cases — speciality steel and advanced semiconductors — the mechanism of mastery loss is identical. An actor stops investing in fundamental research and upmarket progression, in favour of short-term profitability optimisation. Meanwhile, other actors invest over long cycles of 10 to 20 years. At the end of these cycles, competence is irreversibly concentrated elsewhere. Reconquest is not impossible — but it takes the same time as the abandonment did, and costs exponentially more.

◆ WHAT EUROPE HAS INVESTED AND WHAT IT HAS PRODUCED

The European Chips Act of 2023 mobilised €43 billion to attempt to bring Europe's share of global semiconductor production to 20% by 2030. This is a signal. It is not yet a capability. Intel has announced factories in Germany and Poland — projects that take 5 to 10 years to build and reach technical maturity. TSMC has announced a factory in Germany for previous-generation nodes. These investments will reduce dependence on certain segments — they will not recreate the leading edge of technology in Europe within the coming decade.

◆ NASSIHA — THE EUROPEAN CHIPS ACT IS NOT A SHORT-TERM SOLUTION

The European Chips Act is a correct and necessary political decision. It does not correct current dependence — it begins to build the conditions for reducing it over 10 to 15 years. European datacentres running today and tomorrow will continue to depend on chips manufactured outside Europe throughout this entire period. Long-term investment and short-term risk management are two distinct problems that require distinct responses.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

THE H100 GPUs AND THE NEW ARMS RACE

The emergence of generative AI since 2022 has added a further dimension to semiconductor dependence. Large language models, computer vision systems, recommendation engines — all require massively parallel computing GPUs. NVIDIA dominates this market with its H100, H200, and successive generation GPUs. These chips are manufactured by TSMC. Global demand is so intense that it regularly exceeds available production capacity.

For European actors seeking to develop sovereign AI capabilities — whether states, companies or research centres — access to leading-edge GPUs is rationed by TSMC's production capacity and NVIDIA's allocation priorities. This is not a free market. It is a discretionary allocation market in which the best-positioned actors (American hyperscalers, large Asian players) are served first.

◆ AMERICAN EXPORT CONTROLS: A GEOPOLITICAL TOOL

The US government has implemented strict controls on the export of advanced chips to China starting in 2022, progressively extended and tightened in 2023 and 2024. These controls demonstrate a fact that the Opération Dindon corpus has documented from a different angle: the hardware components of digital infrastructure are instruments of foreign policy. A US administration can, by executive order, decide which countries have access to which chips. This power requires no contractual negotiation — it is exercised by unilateral regulatory means.

◆ WHAT THIS MEANS FOR EUROPE

Europe is not the target of current export controls. But it is not immune to a reconfiguration of American geopolitical priorities either. A trade partner can become a strategic competitor. Today's exemptions are not tomorrow's guarantees. A Europe that depends on chips whose access is controlled by a third-party government is a Europe whose digital sovereignty is conditional — conditional on the lasting goodwill of that third-party government.

◆ NASSIHA — THIS IS NOT AN ANTI-AMERICAN ARGUMENT

Europe's dependence on chips manufactured outside its territory is not a problem of American or Taiwanese hostility. It is a problem of risk structure. Unilateral dependence on a critical resource is a structural risk regardless of the quality of current relations with the supplier. The Opération Dindon corpus has made the same demonstration for cloud contracts: it is not that hyperscalers are malevolent — it is that unilateral dependence is structurally fragile.

Amine RAITI · CC BY-NC-SA 4.0
RAITI

5

SECTION 5 · WHAT THE SRE CAN DO

THE OPERATIONAL RESPONSE TO HARDWARE DEPENDENCE

Semiconductor dependence is an industrial policy problem that the SRE cannot solve alone. But the SRE is not without leverage. Operational responses to hardware dependence exist — infrastructure practices that reduce exposure to supply disruption risk without waiting for Europe to have rebuilt its production capacity.

◆ LEVER 1 — EXTENDING HARDWARE LIFETIME

The culture of short-cycle hardware renewal (3 to 5 years) is an industry norm that serves manufacturers' interests more than operators'. A properly maintained server can run for 8 to 10 years on stable workloads. Deliberately extending renewal cycles reduces dependence on the immediate availability of new components — and also reduces the exposure surface to supply disruption risk. This is FinOps applied to hardware, with a strategic dimension that standard FinOps analyses do not integrate.

◆ LEVER 2 — HARDWARE ARCHITECTURE DIVERSIFICATION

The standardisation of datacentre architectures around a few processor families (Intel x86, AMD EPYC, NVIDIA GPU) creates dependence on converging supply chains. The rise of ARM architectures in datacentres — Ampere Computing, AWS Graviton — and open-source RISC-V processors represents genuine diversification. A heterogeneous hardware estate is more resilient to a specific supplier disruption than a homogeneous one.

◆ LEVER 3 — STRATEGIC STOCKS AND LONG-TERM PLANNING

Large organisations — states, telecoms operators, critical infrastructure companies — can integrate into their hardware planning the building of strategic stocks of critical components. This is not common practice in the IT industry, but it is standard in other sectors with critical dependencies (energy, pharmaceuticals, food). The awareness that semiconductors are a strategic resource on a par with oil or medicines has not yet translated into storage practice in the digital sector.

◆ NASSIHA — THESE LEVERS DO NOT SOLVE THE STRUCTURAL PROBLEM

Extending hardware lifetime, diversifying architectures, building stocks — these practices reduce risk exposure but do not eliminate it. They buy time. They do not reconstitute European manufacturing capacity. They are necessary in the immediate term and insufficient over the long term. The structural response — investing in advanced chip manufacturing in Europe — is an industrial policy decision played out at the level of states and decades, not at the level of infrastructure teams and quarters.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

RECOGNISING HARDWARE AS CRITICAL INFRASTRUCTURE

The Opération Dindon corpus has documented dependence on hyperscalers at the level of contracts, certifications and training. One layer was missing below: dependence on the hardware itself. Before being captive to AWS or Azure, one is captive to the chips that run AWS and Azure's servers. Digital sovereignty that stops at contracts and ignores hardware is a façade sovereignty.

◆ MEASURE 1 — INTEGRATE SEMICONDUCTORS INTO INFRASTRUCTURE RISK ANALYSES

Every critical infrastructure risk analysis (business continuity plan, BCP, DRP) should include a "hardware dependence" component documenting the geographic origin of critical components, alternative sourcing lead times, and extended disruption scenarios. This practice is standard for software and contractual risks — it is absent for hardware risks. Adding it costs nothing and makes visible a dependence that is currently ignored.

◆ MEASURE 2 — SUPPORT OPEN ARCHITECTURES SUCH AS RISC-V

The RISC-V instruction set architecture is open source — anyone can design and manufacture a RISC-V processor without paying a licence to Intel or ARM. European research and industrial actors are investing in RISC-V for embedded systems and accelerators. Supporting this ecosystem at the level of public procurement policies and research funding creates, over the long term, an alternative to proprietary architectures concentrated in the United States and the United Kingdom.

◆ MEASURE 3 — TRAIN INFRASTRUCTURE ENGINEERS IN THE HARDWARE LAYER

The Foundation of Iron begins at electricity and hardware (weeks 1 and 6) precisely because understanding the physical layer is the condition for understanding all the layers above. This pedagogy must be extended to the hardware supply chain: where components come from, who makes them, what the alternatives are, how to assess the resilience of a hardware architecture. An SRE who ignores the upstream chain of their infrastructure is an SRE who does not fully understand the risks they are managing.

◆◆◆

No digital progress without mastery of the digital iron. The chip is the ingot of the 21st century. Who makes it decides who gets to build.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

Structural Study · Opération Dindon · June 2026 · Verified Document

THE GUARDIANS OF IRON

While Europe Lays Off Its Bare-Metal Engineers, Hyperscalers Are Training Their Own

RATIO

◆ METHODOLOGICAL WARNING

Every quantified or named claim in this document has been independently verified through direct research — sources: Amazon.jobs, Microsoft Local, Google Data Centers, Indeed. No figure is an unsourced estimate. Wherever information is missing, the document states so explicitly rather than inventing it.

Pages	Hyperscalers	Verified Programmes	Watermark
7	3	3	RATIO DOMINANT

Amine Raiti — Infrastructure Architect & SRE

Former engineering school professor · Trainer since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026 · Claude/Gemini audit, 3 rounds

CC BY-NC-SA 4.0

WHILE THE EUROPEAN MARKET DEVALUES IRON, HYPERSCALERS TRAIN IT IN-HOUSE

Several Opération Dindon corpus studies have already documented the devaluation of physical infrastructure skills in Europe — NoOps, The Dissolution of the Boundary, Iron Mixity. While this movement is confirmed, an inverse, verifiable and dated fact is occurring at the three major American hyperscalers: they are massively recruiting and training, in-house, exactly the skills Europe is abandoning.

◆ THE STARTING POINT — A SAMPLE OF REAL JOB POSTINGS

This study originated from direct examination of Amazon job postings on LinkedIn and Amazon.jobs: Chief Engineer Operations, Controls Engineer EMEA, Critical Infrastructure Mechanical Engineer. All require reading electrical diagrams, configuring industrial controllers (PLC/DCS), HVAC systems maintenance, root cause analysis on equipment failures. These are electrical-trade and automation-engineer job descriptions — not cloud developer ones.

◆ WHAT VERIFICATION REVEALED — A TRIPTYCH, NOT AN ISOLATED CASE

The initial investigation into Amazon led to discovering two structurally identical programmes at Microsoft and Google. This document records all three, with verified sources, and questions what this convergence — not concerted to our knowledge — structurally produces on the European technical training market.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · AMAZON — THE WORK-BASED LEARNING PROGRAM

FROM "NO EXPERIENCE, RIGHT OUT OF HIGH SCHOOL" TO CHIEF ENGINEER

◆ THE PROGRAMME — VERIFIED ON AMAZON.JOBS

The Work-Based Learning Program (WBLP) is a 12-month, paid training programme covering the entirety of datacentre operations — logistics, infrastructure, decommissioning. It combines self-directed web instruction, instructor-led sessions, and on-the-job training. Testimonial cited on the official page: "I came to AWS with no experience, right out of high school. I applied through the WBLP [...] eventually becoming a chief engineer, managing the maintenance activities in multiple facilities."

◆ CONFIRMED EUROPEAN PRESENCE

A "Data Center Chief Engineer, DCEO" position is explicitly located in the "Greater Milan Area" (Job ID 3111099, Amazon.jobs). The exact technical vocabulary found across multiple different job postings: BMS (Building Management System), EPMS (Electrical Power Monitoring System), industrial protocols MQTT, BACnet, MODBUS. This is the exact language of an industrial automation engineer.

◆ THE NECESSARY NUANCE — CAPTIVE BY THE SYSTEM, NOT BY THE CONTRACT

The WBLP trains generic electrical and automation skills, theoretically transferable to any industrial employer. What makes these profiles captive is not a non-compete clause, but the specificity of the proprietary internal systems they are trained on (Amazon BMS/EPMS) — code and architecture lock more effectively than law.

LEX

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · MICROSOFT AND GOOGLE — THE SAME MOVEMENT, INDEPENDENTLY

THREE HYPERSCALERS, THREE PROGRAMMES, THE SAME LOGIC

◆ MICROSOFT DATACENTER ACADEMY — 11 GLOBAL SITES, 3 NAMED IN EUROPE

Verified on Microsoft Local and Microsoft Careers: 11 global locations, including Dublin (Ireland), Amsterdam/Hoorn/Schagen (Netherlands), Gävle/Sandviken (Sweden), and a 2025 launch in Finland (Vihti, Kirkkonummi, Espoo) in partnership with vocational schools Omnia and Luksia. Direct quote from Mervi Airaksinen, Microsoft Finland General Manager: the programme trains in "ICT skills, building mechanics, electricity and automation." Microsoft supplies lab equipment (racks, servers, wiring) and funds scholarships.

◆ GOOGLE — STAR PROGRAM AND EUROPEAN APPRENTICESHIP CONFIRMED

A "Data Center Technician Apprentice" position is open in the Netherlands for a January 2026 start (Google Careers, directly verified). The Skilled Trades and Readiness (STAR) Program, documented on Google Data Centers, received a $50 million investment in 2026 to train over 300,000 Americans in electrical, plumbing and sheet metal trades — physical infrastructure, not cloud. Google's apprenticeship programme (18 months, paid) is confirmed available in France, Ireland, Switzerland and the Netherlands.

◆ THE TRIPTYCH — THREE STRUCTURALLY IDENTICAL STRATEGIES

WBLP (Amazon), Datacenter Academy (Microsoft), STAR Program (Google) share the same architecture: training profiles with no prior experience, supplying pedagogical equipment, funding scholarships, and direct pathway into employment at the hyperscaler. No evidence of coordination between the three actors was found — these are three independent responses to the same economic constraint: the European technical labour market no longer produces enough Iron-trained profiles.

LEX

Amine RAITI · CC BY-NC-SA 4.0

SKILLS DO NOT DISAPPEAR — THEY CHANGE OWNER AND JURISDICTION

◆ THE TCO CALCULATION THAT MASKS THE REAL TRANSFER

The European market justifies the devaluation of infrastructure skills with a classic TCO calculation: outsourcing servers and networks turns fixed costs (CapEx) into variable costs (OpEx). The framing error of CIOs is believing this transfer of responsibility makes the need for physical skills disappear. It does not disappear — it changes owner.

◆ THE INDUSTRIALISATION ASYMMETRY

European CIOs destroy their infrastructure skills because they lack the critical scale to make them profitable at a single company's level. Hyperscalers, conversely, build economies of scale that make total internalisation of Iron maintenance profitable — and necessary, since a single outage at their scale affects thousands of customers simultaneously.

◆ THE DIAGNOSTIC ASYMMETRY DURING OUTAGES

The job postings analysed require mastery of root cause analysis on industrial equipment failures. In a typical European company, this analysis is largely delegated to the equipment manufacturer or an external subcontractor. By internalising this skill, hyperscalers accumulate operational intelligence that European companies have largely abandoned. Practical consequence: during a major infrastructure failure, the European client company no longer has the in-house expertise to understand the root cause — it depends on the hyperscaler's diagnosis.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THREE INDEPENDENT ACTORS, THE SAME EFFECT ON THE SAME SCHOOL CATCHMENT AREAS

◆ WHAT IS OBSERVABLE — NOT COLLUSION, A STRUCTURAL CONVERGENCE

Examined individually, the WBLP, the Datacenter Academy and the STAR Program appear as healthy recruitment competition. But their geographic location converges on the same technical and vocational public school catchment areas — the Netherlands, Ireland, Finland notably appear in several of the three programmes. No evidence of coordination between the three hyperscalers was found, and it would be legally inaccurate to speak of a cartel without proof of collusion — what is documented is an independent economic logic producing the same structural effect.

◆ WHAT THIS CONVERGENCE STRUCTURALLY PRODUCES

The three hyperscalers do not merely recruit: they supply pedagogical equipment (racks, servers, wiring), fund scholarships, and by construction steer the practical content of training toward their own systems. An electrician or automation engineer trained in these pathways learns on equipment and protocols chosen by the hyperscaler — not necessarily what maintaining a classic factory, a local SME, or an independent sovereign infrastructure would require.

◆ THE LIMIT OF THIS ARGUMENT — WHAT IS NOT PROVEN

This study does not claim public curricula are "dictated" by hyperscalers in the sense of total control — the programmes consulted speak of "guidance" and "curriculum alignment," not imposition. Claiming a graduate of these pathways "will no longer be able" to maintain an independent infrastructure would be an unverified extrapolation. What is documentable: a strong material and pedagogical orientation toward the proprietary systems of the same three actors, in the same territories.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

IN CLOSING

NOT A DELIBERATE SUICIDE — AN UNCONSCIOUS ONE

No French company laying off a Bare-Metal engineer does so thinking "this way, Amazon will recruit better than us." Each individual devaluation decision follows its own short-term TCO logic. But the cumulative effect, documented here across three independently verified actors, is unambiguous: while Europe sheds its physical infrastructure skills, the three American hyperscalers recruit them, train them, and internalise them — all the way down to public school.

◆ WHAT IRON NEVER LIES ABOUT

A Chief Engineer trained by the WBLP, an automation engineer graduated from the Datacenter Academy, a technician certified by the STAR Program — all master a skill Europe documents as obsolete at the very moment three of the most highly valued companies in the world invest tens of millions of dollars to pass it on to candidates with no prior experience. Iron does not lie. The investments confirm it.

Nobody decided to lose mastery of Iron.

A thousand individually rational decisions led there.

Three foreign actors, without coordinating, drew the same conclusion.

Amine RAITI · *Opération Dindon* · 2026

NEMO SUPRA LEGEM EST

CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE DIGITAL GROUND

Where the datacentres are, who decides, who pays

◆ CONTEXT

The Opération Dindon corpus covered contractual dependence on hyperscalers, semiconductor dependence, training and competence recognition. It had not addressed the physical and geographical dimension: where European datacentres are, who owns them, what legal constraints apply to the data they hold. This study fills that blind spot.

Amine Raiti · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE MAP

AMSTERDAM, FRANKFURT, DUBLIN — THREE HUBS, ONE DEPENDENCE

The geography of datacentres in Europe is not neutral. It results from fiscal, energy and political decisions that have directed investment towards certain territories. Three zones concentrate the bulk of capacity: the Amsterdam-Frankfurt corridor (AMS-IX and DE-CIX internet exchange nodes), Ireland (European headquarters and datacentres of American hyperscalers), and the Nordic countries (abundant energy, natural cooling, attractive tax regimes).

◆ WHY IRELAND BECAME EUROPE'S DATACENTRE

Ireland hosts the European datacentres of Google, Meta, Amazon, Microsoft and Apple. Not for any particular technical quality of its soil — but for its 12.5% corporate tax rate, which attracted the European headquarters of hyperscalers, bringing their datacentres in their wake. Data from hundreds of millions of Europeans thus flows through infrastructure physically located in Ireland, subject to the European GDPR, but operated by American companies subject to the CLOUD Act.

◆ THE CLOUD ACT: THE LAW THAT FOLLOWS DATA

The CLOUD Act (2018) authorises US authorities to demand access to data held by American companies, regardless of where it is physically stored. An AWS datacentre in Dublin hosting a French SME's data is simultaneously subject to GDPR and the CLOUD Act. In case of conflict between the two, the legal answer remains uncertain. This uncertainty is structural, not accidental.

◆ NASSIHA — WHAT GDPR DOES NOT RESOLVE

GDPR compliance does not immunise data against a US DOJ request based on the CLOUD Act. Both regimes apply simultaneously — and their interaction is not resolved by current international law.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · OWNERSHIP, OPERATION, CONTROL

THREE LEVELS THAT DO NOT ALWAYS COINCIDE

The question "who owns European datacentres" breaks down into three distinct levels: ownership of the building and land, technical operation of the infrastructure, and control of the data flowing through it. A datacentre can be physically on European soil, operated by a European subsidiary, and yet subject to extraterritorial control by an American parent company.

◆ SOVEREIGN EUROPEAN ACTORS: A SHORT LIST

OVHcloud, Hetzner, Scaleway, Data4 — actors who both own and operate large-scale datacentres under European control are few. The vast majority of public cloud capacity in Europe is operated by AWS, Azure and GCP — American legal entities, subject to the CLOUD Act, whatever the physical location of their servers.

◆ THE INTERXION ACQUISITION BY DIGITAL REALTY — $8.4 BILLION (2020)

Interxion was the leading independent European neutral datacentre operator (colocation). Its 2020 acquisition transferred ownership and control of critical European infrastructure to an American company. An ordinary commercial transaction — but one that illustrates the mechanism by which European digital infrastructure sovereignty erodes not through conquest, but through acquisition.

◆ THE SAME MECHANISM AS VOLVO

"The Mastery of Iron" documented how Volvo was acquired by Geely — not through a hostile act, but through the absence of an industrial policy. Europe does not lose its datacentres through expropriation. It cedes them through a lack of strategic digital asset protection policy.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · ENERGY AS A STRATEGIC VARIABLE

THE DATACENTRE IS AN INVERTED POWER STATION

A medium-sized datacentre consumes as much electricity as a city of 50,000 inhabitants. This consumption is both a siting constraint and a political lever. The Nordic countries won the hyperscaler datacentre competition precisely because they built attractiveness policies around energy: abundant, cheap, low-carbon hydro and nuclear power; low ambient temperatures allowing natural free cooling.

◆ FRANCE: AN AMBIGUOUS POSITION

France has real assets — low-carbon nuclear electricity, dense fibre, a market of 67 million people. But it has not developed a coherent policy of attractiveness and protection simultaneously. It attracted hyperscaler datacentres (AWS Paris region) without demanding in return any commitments on public data localisation, technology partnerships with French actors, or contributions to fibre infrastructure financing.

◆ SUBMARINE CABLES: THE INVISIBLE INFRASTRUCTURE

Datacentres are connected to each other and to the rest of the world by submarine fibre optic cables carrying 99% of international internet traffic. Since 2010, American hyperscalers own or co-finance a growing share of this capacity. In 2022-2023, several cables were severed in the Mediterranean and Baltic in unexplained circumstances — revealing a real, exploitable physical vulnerability that is nearly impossible to protect across thousands of kilometres of ocean.

◆ NASSIHA — SOVEREIGNTY IS NOT AUTARKY

Preserving sovereignty over data and infrastructure does not mean expelling hyperscalers. It means ensuring dependence remains reversible. Reversibility is the true metric of sovereignty — not the absence of dependence, but the absence of irreversible dependence.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

A TENSION EUROPE HAS NOT RESOLVED

Europe faces a real tension between attracting hyperscaler investment and preserving sovereignty over its data and infrastructure. These two objectives are not incompatible — but articulating them requires an explicit policy that most Member States have not formulated. The intra-European competition to attract datacentres pushes each country to offer more without demanding anything in return — a negative-sum game for collective sovereignty.

◆ GAIA-X: THE DILUTED AMBITION

GAIA-X, launched in 2019 as a European sovereign cloud initiative, was progressively diluted by the integration of hyperscaler members — AWS, Azure and GCP are GAIA-X members and influenced its standards in a direction favourable to their interests. Its partial failure is instructive: digital sovereignty is not built by including the actors one seeks to emancipate from.

◆ SITING CONDITIONS AS A NEGOTIATION LEVER

A state hosting a hyperscaler datacentre has a unique negotiation moment — before the building permit, before the electrical connection, before the tax benefits. This moment is rarely used to demand sovereignty counterparts: EU localisation of public data, training of local technicians, third-party access for European actors, non-unilateral application of extraterritorial legislation to European data. These counterparts are negotiable. They are neglected because states compete with each other.

◆ NASSIHA — INTRA-EUROPEAN COMPETITION IS THE REAL PROBLEM

As long as Ireland can attract AWS datacentres by offering 12.5% corporate tax and France cannot oppose this without a European agreement, fiscal and regulatory European coordination is the prerequisite for any coherent digital sovereignty policy. Without it, each state plays alone against hyperscalers who play collectively on a global scale.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

GEOGRAPHY IN ARCHITECTURE DESIGN

The geography of datacentres is not a geopolitical abstraction reserved for policymakers. It is an architectural constraint that every SRE must integrate into their deployment decisions. Where is the data? Under which jurisdiction? What are the latency times between regions? What are the connectivity disruption scenarios? These are not optional questions — they are structural.

◆ LEVER 1 — KNOW THE LEGAL ORIGIN OF INFRASTRUCTURE USED

Every SRE deploying on public cloud should know: what is the legal entity operating this cloud in the relevant country, is it subject to the CLOUD Act, what are the data localisation clauses in the service contract. This information should appear in every infrastructure risk analysis — alongside the availability SLA or backup policy.

◆ LEVER 2 — DESIGN FOR GEOGRAPHIC RESILIENCE

A resilient architecture does not only distribute load between availability zones (AZ) of the same provider in the same country. It considers resilience across distinct providers, distinct countries, and distinct legal regimes. Multi-cloud and multi-jurisdictional are not the same thing — but the combination of the two is true geographic resilience.

◆ LEVER 3 — INTEGRATE SUBMARINE CABLES INTO CONTINUITY PLANS

A BCP/DRP that does not model the submarine cable severance scenario affecting connectivity between regions is an incomplete plan. These incidents are documented, recurring, and slow to repair. The architectural response — backup routing, local cache, graceful degradation — must be prepared in advance, not discovered during the incident.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

THE DIGITAL GROUND AS A STRATEGIC ASSET

The proposal of this study is to treat the digital ground — the physical location of datacentres, the cables connecting them, the energy powering them — as a strategic asset on a par with agricultural land or the railway network. This is not nationalisation. It is the recognition that the physical infrastructure of digital is a critical infrastructure that deserves explicit governance.

◆ MEASURE 1 — MAP AND PUBLISH GEOGRAPHIC DEPENDENCE

Each Member State should publish an annual inventory of datacentre capacity on its territory: owner, operator, volume, energy consumed, legal origin of controlling entities. This transparency does not exist today. It would allow citizens, regulators and businesses to assess their real exposure to extraterritorial legal risk.

◆ MEASURE 2 — CONDITION TAX BENEFITS ON SOVEREIGNTY COMMITMENTS

Any tax benefit or administrative facilitation granted to a hyperscaler datacentre should be conditional on verifiable commitments: EU localisation of public data, training of local technicians, third-party access for European actors, and non-unilateral application of the CLOUD Act to European data hosted in these datacentres.

◆ MEASURE 3 — PROTECT SOVEREIGN EUROPEAN ACTORS

OVHcloud, Hetzner, Scaleway and their European equivalents deserve explicit preference in public procurement — not out of protectionism, but for consistency with a stated digital sovereignty policy. A state that preaches digital sovereignty while hosting its data on AWS is not practising sovereignty — it is talking about it.

◆◆◆

The digital ground is not an abstraction. It is concrete, cable, electricity and law. Whoever controls these four elements controls the infrastructure. Not the contract. The ground.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE TAIWAN BOTTLENECK

TSMC · ASML · The Strait as Single Point of Failure of the Global Digital Economy

RATIO

◆ THE THESIS

Hyperscaler lock-in has three layers. The contractual layer — documented in Terms Under the Microscope. The software layer — documented in The Newspeak That Costs Dear. The material layer — the deepest, least visible, the one the Terms do not mention because it precedes the Terms. This study documents the first dimension of the material layer: geopolitical dependency on the Taiwan Strait — and why the organisation that owns its physical hardware is the only one with business continuity independent of this dependency.

TSMC SUB-5NM	EUV MONOPOLY	WATERMARK
~90%	**ASML 100%**	**RATIO**

Amine Raiti — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · THE PHYSICAL MONOPOLY OF SILICON

DESIGNING A FREE CHIP IS POINTLESS IF YOU CANNOT FAB IT

The Opération Dindon corpus documented RISC-V as the open-source ISA alternative to Intel, AMD and ARM proprietary architectures. This documentation is correct and necessary. It is insufficient. RISC-V is an instruction set architecture — a logical design. For a RISC-V chip to exist physically, it must be fabricated in silicon. And fine-node silicon fabrication is subject to two inescapable physical monopolies.

◆ MONOPOLY I — ASML AND EUV MACHINES

Extreme ultraviolet lithography (EUV) is the fabrication technique that enables chips below 7nm — the nodes used in datacentre GPUs, smartphone processors, and AI chips. ASML, a Dutch company based in Eindhoven, is the world's sole manufacturer of EUV machines. One ASML EUV machine costs approximately €150 to €200 million, weighs 180 tonnes, requires 40 containers to transport, and is assembled from 100,000 components from 5,000 suppliers in 30 countries.

There is no alternative. Not in China (despite massive investment since 2019), not in Russia, not in Europe other than ASML itself. Any fab that wants to produce advanced chips must go through ASML. And ASML, under pressure from the US and Dutch governments, has not been selling its EUV machines to China since 2023.

◆ MONOPOLY 2 — TSMC AND FINE-NODE FABRICATION

TSMC (Taiwan Semiconductor Manufacturing Company) produces approximately 90% of the world's advanced chips (nodes below 5nm). Its customers include Apple, NVIDIA, AMD, Qualcomm, Google, Amazon, Microsoft. The NVIDIA H100 GPUs powering AI models — including those of Anthropic, OpenAI and Google DeepMind — are fabricated by TSMC in Taiwan.

TSMC in Taiwan. On an island of 36,000 km² separated from mainland China by a 130 km strait. In a zone subject to territorial claims by the People's Republic of China. In a region where large-scale Chinese military exercises took place in 2022 and 2023.

◆ WHAT THIS MEANS FOR THE CIO WHO MIGRATED 100% TO CLOUD

Their cloud vendor (AWS, GCP, Azure) is itself dependent on TSMC for new GPU servers. In a Taiwan crisis, hyperscalers will ration new capacity — first for their own internal needs, then for customers in contractual priority order. The CIO whose business continuity depends on continuous availability of cloud resources does not control this rationing.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE BOTTLENECK FIGURES

WHAT A 6-MONTH TSMC INTERRUPTION WOULD PRODUCE — MEASURED

INDICATOR	FIGURE	SOURCE / NOTE
TSMC share of world advanced chips <5nm	**~90%**	SemiAnalysis 2023
Global EUV machine manufacturers	**1 (ASML)**	Absolute global monopoly
Estimated economic impact of 6-month TSMC disruption	**$600Bn–$1,200Bn**	PIIE 2023 · probable global recession
Time to build equivalent fab from scratch	**10 to 15 years**	US CHIPS Act · Intel German fab difficulties 2024
NVIDIA H100 GPU (core of AI datacentres) — manufacturer	**TSMC (Taiwan)**	AWS, GCP, Azure depend on this GPU
US→China advanced chip export restrictions (in force)	**Since Oct. 2022**	Precedent of technological extraterritoriality

◆ WHAT US→CHINA EXPORT RESTRICTIONS REVEAL FOR EUROPE

Since October 2022, the US government has banned the export to China of advanced chips and associated production equipment. This unilateral decision was applied extraterritorially — forcing ASML (a Dutch company) to cease deliveries to China. It was applied without consulting the European Union.

This precedent says something important for European organisations using US cloud: the United States has demonstrated its capacity and willingness to exercise extraterritorial control over the silicon supply chain. If geopolitical circumstances required it, similar control could be exercised over European access to US cloud services — or over European hardware equipment dependent on US technology.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · REFURBISHED AS GEOPOLITICAL INSURANCE

A DELL POWEREDGE R750 BOUGHT TODAY WILL BE OPERATIONAL IN 5 YEARS — REGARDLESS OF TAIWAN

◆ WHY REFURBISHED IS GEOPOLITICAL INSURANCE

A Refurbished Grade A Dell PowerEdge R750 bought today contains chips fabricated 2 to 3 years ago. These chips already exist in the physical world. Their availability no longer depends on future TSMC production, US export restrictions, ASML machine availability, or the geopolitical situation in the Taiwan Strait. They are there. They work. They will continue to work for 5 to 7 additional years with standard maintenance.

In a silicon supply crisis — geopolitical (Taiwan conflict), natural (earthquake, typhoon in production zones), or regulatory (new export restrictions) — the patrimonial and strategic value of existing physical servers increases. Demand for qualified certified second-hand hardware rises. The organisation that owns its physical hardware has business continuity that the 100% cloud organisation does not.

◆ WHAT HAPPENS TO CLOUD IN A TAIWAN CRISIS

Hyperscalers have server stocks — but limited. In an extended TSMC production interruption, they cannot order new H100 GPUs or new datacentre CPUs. Available capacity will be maintained — but not increased. Customers who want to scale up during a global crisis will hit rationed capacity.

Who gets served first? The hyperscalers themselves (their own AI services, critical internal services), then customers in contractual priority order — Enterprise Agreements first, Pay-as-you-go last. The organisation without an Enterprise contract is last in the queue.

The organisation with its own physical server is outside the queue. It does not wait.

◆ THE FORMULA LINKING THE TWO STUDIES

Refurbished Grade A is 11× cheaper than the GCP VM over 5 years (FinOps study). It is also geopolitical insurance whose value increases exactly when cloud becomes most vulnerable (geopolitical study). These two properties are not independent — they are two faces of the same decision: own your hardware or rent someone else's.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · WHAT THE CIO MUST CALCULATE — NOT IN THE STANDARD TCO

THE COST OF GEOPOLITICAL DEPENDENCY IS NOT IN THE AWS INVOICE

◆ WHAT THE STANDARD TCO FORGETS

The standard Total Cost of Ownership calculation cloud vs on-premise compares: hardware acquisition cost + staff + maintenance + datacentre vs monthly cloud invoice. This is a correct calculation in a stable world. It is insufficient in a world where the silicon supply chain is concentrated on a geopolitically disputed island.

What the standard TCO does not include:

— The geopolitical insurance premium of existing physical hardware

— The cost of cloud rationing in a supply crisis

— The value of business continuity independent of the global supply chain

— The cascade dependency: cloud → H100 GPU → TSMC → Taiwan Strait

◆ THE THREE LAYERS OF HYPERSCALER LOCK-IN

Layer 1 — Contractual: Terms Section 14.12 (California jurisdiction), noncancellable commits, egress fees, unilateral price modification. Documented in Terms Under the Microscope.

Layer 2 — Software: proprietary dialects (DynamoDB, BigQuery, Lambda), proprietary certifications, newspeak that destroys neutral skills. Documented in The Newspeak That Costs Dear.

Layer 3 — Material: dependency on TSMC for new GPU servers, dependency on ASML for EUV machines, geopolitical concentration in the Taiwan Strait. This layer is the deepest. It is also the only one the Terms do not mention — because it precedes the Terms. The cloud vendor does not guarantee its own silicon supply chain continuity. Nobody guarantees it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · EUROPE AND SILICON — WHAT EXISTS AND WHAT IS MISSING

RISC-V · INTEL GERMAN FAB · EUROPEAN CHIPS ACT · THE HORIZON AND THE REALITY

◆ WHAT EXISTS — EUROPEAN INITIATIVES

European Chips Act (2023): €43 billion to double Europe's share of global semiconductor production — from 10% to 20% by 2030. Ambitious objective. Fine-node fabrication capacity: not yet achieved.

Intel Fab in Germany (Magdeburg): planned €33 billion investment, 1.8nm fabrication announced. Delays: pushed back in 2024 due to reduced subsidies and market difficulties. Planned capacity: 2030 at earliest.

RISC-V: open-source instruction set architecture supported by an international foundation. RISC-V chips exist and are produced — primarily by TSMC for customers who order them. The architecture is free. The fabrication remains dependent on the same foundries.

◆ WHAT IS MISSING — THE 10-YEAR GAP

There is no fab in Europe today capable of fabricating chips below 7nm. TSMC and Samsung are the only ones worldwide able to do so at industrial volume. China is investing massively but remains blocked below 28nm due to ASML machine restrictions.

The gap is 10 to 15 years in optimistic scenarios — and assumes announced investments materialise, specialist engineers are trained, ASML machines are available (ASML is itself subject to US pressure on its exports), and markets are there.

The corpus conclusion: RISC-V is the correct horizon. The sovereign foundry is the necessary horizon. Between both and today, there is a 10 to 15-year vulnerability window during which dependency on the Taiwan Strait is structural and unavoidable. It is in this window that the Opération Dindon corpus operates. It is in this window that the CIO's decisions are made.

The CIO who signs for managed cloud accepts three layers of dependency.
The first is in the Terms they signed.
The second is in their developers' code.
The third is in the Taiwan Strait.
This third layer appears in no contract.
It exists nonetheless.

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE MATERIAL BLACK BOX

Intel Management Engine · AMD Platform Security Processor

Lock-In at the Silicon Level

◆ **THE THESIS**

The organisation that believes it has freed itself from cloud by buying bare-metal but runs on Intel or AMD silicon with unauditable proprietary firmware has solved a software dependency problem to fall into a potentially deeper hardware dependency problem. Hyperscaler lock-in is at the software level. Hardware lock-in is at the silicon level. This study documents the second dimension of the material layer: the hidden processors in every server, their nature, and what they mean for sovereignty.

INTEL ME	AMD PSP	WATERMARK
Since 2008	**Since 2013**	**RATIO**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

RATIO

I

SECTION I · THE INTEL MANAGEMENT ENGINE — THE SECOND COMPUTER IN THE SERVER

EVERY INTEL PROCESSOR CONTAINS AN INDEPENDENT COMPUTER LINUX DOES NOT CONTROL

Since 2008, every Intel processor integrates a subsystem called the Intel Management Engine (IME). This is not an optional feature. Not an external component. It is a secondary processor etched into the main CPU die, with its own firmware, its own operating system (MINIX 3), and its own access to memory, network and storage.

◆ WHAT THE INTEL MANAGEMENT ENGINE DOES

The IME was designed for remote enterprise server management — allowing a sysadmin to restart, reconfigure or diagnose a server without being physically present. This function is legitimate and useful. What is problematic is its architecture:

It runs independently of the OS: Linux, Windows, BSD — no operating system controls the IME. It runs below the kernel level. It is invisible to standard audit tools.

It functions when the server is "off": as long as the server is powered (even in deep standby), the IME is active. It can receive network instructions and execute them without the main OS being started.

It has independent network access: the IME can communicate on the network via channels distinct from those used by the OS. On some motherboards it has its own dedicated network port.

Its firmware is proprietary: Intel does not publish the IME firmware source code. It cannot be audited or easily replaced. Partial disabling attempts exist (me_cleaner) but are not guaranteed complete.

◆ AMD PLATFORM SECURITY PROCESSOR — THE SAME PROBLEM ON AMD SIDE

Since 2013, AMD processors integrate the Platform Security Processor (PSP) — also called AMD Secure Technology. It is based on an ARM Cortex-A5 architecture, has proprietary firmware, and similar functions to the IME: pre-OS execution, memory access, hardware-level security management.

AMD has published more technical specifications than Intel on this component, but the firmware remains proprietary and not fully auditable. The coreboot community is working on alternative implementations but coverage is partial.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

BUYING BARE-METAL WITHOUT MASTERING IME/PSP — RELOCATING THE LOCK-IN

◆ THE FIRMWARE JURISDICTION CHAIN

The Intel IME firmware is developed and distributed by Intel Corporation, headquartered in Santa Clara, California. Firmware updates are cryptographically signed by Intel. They are distributed via the motherboard manufacturer's update channels (Dell, HP, Lenovo, Supermicro) which incorporate them into their own BIOS/UEFI updates.

This firmware is subject to US law. It is subject to National Security Letters. It is subject to the CLOUD Act in its dimension of information system control. An organisation running sensitive data on Intel servers believing it has freed itself from US cloud is in fact running that data on unauditable US firmware.

This is not speculation. It is the logical consequence of IME architecture and US law as documented in Terms Under the Microscope.

◆ WHAT THE SRE CANNOT DO WITH THE IME

An experienced Linux administrator has access to everything: kernel, drivers, logs, hardware metrics, system calls. They can audit virtually any layer of the system. The IME is the exception — the layer that escapes their technical authority.

They cannot read IME logs without special (and partial) tools. They cannot inspect IME network traffic with tcpdump — the IME operates below the network interface. They cannot disable the IME without risking destabilising the server (some BIOS functions depend on IME). They cannot verify IME firmware integrity with standard audit tools.

The Technical Primacy documented in the corpus stops at the IME's surface. Below: proprietary terra incognita.

◆ THE PARTIAL ANSWERS THAT EXIST

Coreboot: open-source alternative firmware for BIOS/UEFI, supported on a limited number of motherboards. Can reduce IME exposure without eliminating it completely.

ME_CLEANER: tool attempting to neutralise non-essential IME modules. Partial results, destabilisation risks, not officially supported.

RISC-V + open processors: the 10-15 year horizon. RISC-V processors without IME exist but are not yet available in high-performance datacentre versions.

IBM POWER + OpenPOWER: existing alternative with a higher level of auditability than x86 processors. Different cost and software ecosystem.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE THREE LOCK-IN LAYERS — COMPLETE SYNTHESIS

CONTRACTUAL · SOFTWARE · MATERIAL — TOTAL DEPENDENCY MAPPED

LAYER I — CONTRACTUAL

What it is: California jurisdiction (GCP §14.12), noncancellable commits, egress fees, unilateral price modification, possible government access (CLOUD Act).
How to exit: change vendor, renegotiate contracts, choose sovereign MSP. Difficult but possible — documented in The Accompanied Exodus.
Corpus study: Terms Under the Microscope.

LAYER 2 — SOFTWARE

What it is: proprietary dialects (DynamoDB, BigQuery, Lambda), proprietary certifications, newspeak destroying neutral skills, application code speaking the vendor's language.
How to exit: Sovereign Interface (Anti-Corruption Layer), progressive migration (Gentle Exit), rebuilding neutral skills (Foundation of Iron). Slow but doable — documented in the corpus.
Corpus studies: Sovereign Interface · The Newspeak That Costs Dear.

LAYER 3 — MATERIAL

What it is: dependency on TSMC/ASML for new GPU servers (The Taiwan Bottleneck), proprietary unauditable Intel ME / AMD PSP firmware (The Material Black Box), strategic value of existing hardware in supply crisis (The Strategic Refurbished).
How to exit: own existing hardware (immediate horizon), coreboot/me_cleaner (partial, complex horizon), RISC-V + sovereign foundries (10-15 year horizon). No simple solution today.
Corpus studies: Digital Iron · Taiwan Bottleneck · Material Black Box · Strategic Refurbished.

RATIO

The contractual layer is visible — it is in the Terms.
The software layer is visible — it is in the code.
The material layer is invisible — it is in the silicon.
The Opération Dindon corpus documented all three.
The third layer is the only one the Terms do not mention.
It exists nonetheless.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE FINOPS OF THE LAST GRAM

Precious metal recovery

FROM END-OF-LIFE BARE-METAL HARDWARE

DIY Lab Amine RAITI — From rack to crucible · SysAdmin Alchemy

◆ **CONTEXT OF THE STUDY**

After every rack/derack operation, end-of-life servers accumulate. The traditional WEEE recycling chain pays between €0 and €50 per server — or charges for collection. This study documents a radical alternative: direct recovery of precious metals contained in obsolete hardware. Gold, copper, silver, palladium — the residual value is real, measurable, and extractable with accessible lab equipment.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE ECONOMIC THESIS — WHAT A SERVER IS REALLY WORTH

THE RESIDUAL VALUE NOBODY CALCULATES

A 5-year-old Dell R640 server is considered an end-of-life cost. It costs money to recycle, store, move. Nobody in IT calculates what it contains in gross metallurgical value. This calculation exists — and it changes the perspective.

◆ METAL CONTENT OF A STANDARD 1U SERVER

Gold (Au): 0.2 to 0.5g depending on model. Present in DIMM gold fingers, PCIe connectors, CPU socket contacts, high-end SAS/SATA connectors, chipset chips. June 2026 rate: ~€85/g. Value: **€17 to €42 per server.**

Copper (Cu): 0.5 to 1.5kg in cables, connectors, thick PCB traces. Rate: ~€9/kg. Value: **€4 to €13 per server.**

Silver (Ag): traces in some contacts and capacitors. 0.1 to 0.3g. Rate: ~€0.95/g. Value: **€0.10 to €0.30 per server.**

Palladium (Pd): present in MLCC capacitors on motherboards and high-end network cards. 0.01 to 0.05g. Rate: ~€35/g. Value: **€0.35 to €1.75 per server.**

Estimated total per R640 server: €22 to €58 in raw precious metals.

For a rack of 40 servers: €880 to €2,300 in recoverable residual value.

◆ COMPARISON WITH THE WEEE CHAIN

The licensed WEEE chain buys servers in bulk — between €0 and €50 per unit depending on condition. It sometimes charges for collection on batches under 500kg. It absorbs the metal value in its margins. Direct recovery bypasses this intermediation — at the cost of a time and lab equipment investment that this study documents and costs out.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

FROM COMPONENT TO CRUCIBLE — TWO OPTIONS

Chemical preparation · DIY Lab Amine RAITI

⚠️ MAJOR WARNING — NON-NEGOTIABLE SAFETY

Handling hazardous chemicals (strong acids, powerful oxidants). **MANDATORY PPE:** acid-resistant gloves, full-seal goggles, minimum FFP3 mask. **VENTILATION ESSENTIAL:** outdoors or chemical fume hood. **NEVER ALONE.** Keep sodium bicarbonate on hand to neutralise splashes. Know the Poison Control Centre number.

◆ OPTION A — PCBs AND CONNECTORS (ACID METHOD)

COPPER DISSOLUTION: immerse connectors in a mixture of hydrochloric acid (HCl 33%) and hydrogen peroxide H_2O_2 (12%). Accessible version: concentrated white vinegar + sea salt + H_2O_2 (12%). Adding an aquarium air pump accelerates copper dissolution and detachment of gold foils — these float to the surface.

PURIFICATION: collect gold foils. Dissolve in 200ml HCl + 2 tbsp sodium nitrate ($NaNO_3$). Heat to 80°C for 15 minutes. Alternative: HCl + a few ml of bleach (⚠️ extremely dangerous — chlorine gas release).

FILTRATION AND PRECIPITATION: filter the solution. Neutralise with sulfamic acid (except bleach method). Precipitate pure gold with iron(II) sulfate ($FeSO_4$). Rest 24h. Brown gold powder settles at the bottom.

MELTING: collect the powder, dry it, melt with torch in graphite or refractory clay crucible.

◆ OPTION B — CHIPS AND COMPONENTS (INCINERATION + CHEMISTRY)

PRE-INCINERATION: separate chips from PCBs. Burn outdoors (barbecue or fireplace) to eliminate epoxy resins and organic carbon. The ash contains concentrated metals.

CHEMICAL TREATMENT: the dissolution and precipitation process is identical to Option A. The yield is slightly lower but gold per gram of treated material is higher.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · NEUTRALISATION — THE STEP THAT CANNOT BE SKIPPED

CHEMICAL WASTE — MANDATORY PROTOCOL

Chemical waste management after extraction · DIY Lab Amine RAITI

⚠️ CRITICAL POINT — NEVER DISCARD SPENT SOLUTIONS AS-IS

Post-extraction solutions contain strong acids and dissolved metals — copper, iron, tin, lead. Pouring them down the drain or into nature is an environmental offence. It is also a beginner's mistake. The neutralisation protocol is as important as the extraction protocol.

◆ CAUSTIC SODA NEUTRALISATION — STEP BY STEP

Materials: caustic soda (sodium hydroxide NaOH, granules or solution), pH paper or pH meter, HDPE plastic container, stirrer.

Protocol:

1. Slowly pour caustic soda into the spent acid solution (never the reverse — risk of violent splashing). The reaction is exothermic — the mixture heats up.
2. Observe the pH rising. Dissolved metals progressively precipitate as coloured sludge — blue-green for copper, rust-brown for iron, grey for tin.
3. Continue until neutral pH (7). Check with pH paper or meter.
4. Leave to settle 24 to 48 hours. Sludge settles at the bottom.
5. Filter the solid sludge. Package in sealed bags labelled "metallic waste — specialist treatment".
6. Drop sludge at a specialist hazardous waste facility (licensed WEEE processor).
7. The neutralised liquid at ~pH 7 can be disposed of in the drainage system — check local regulations.

The copper in the sludge can itself be sold to specialist smelters (Aurubis, Comet Traitements) if volumes are sufficient.

◆ SLUDGE AS A SECOND VALUE SOURCE

Neutralisation sludge from treating 40 servers contains 50 to 200g of precipitated copper, plus traces of nickel and tin. Some recovery smelters (Comet Traitements in Belgium, Umicore) buy this sludge at variable rates by concentration. It is not a fortune — but it is an additional value stream on waste that would otherwise cost money to eliminate.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE PROFITABILITY CALCULATION

THE DIY LAB ROI

The crucible after melting — recovered gold · DIY Lab Amine RAITI

◆ INITIAL INVESTMENT — THE DIY LAB

Safety equipment: acid gloves (€15), full-seal goggles (€25), FFP3 masks x5 (€15), chemical apron (€20). Total: ~€75

Lab equipment: aquarium air pump (€12), HDPE 5L containers x3 (€25), filter funnel + filter paper (€10), pH meter (€15), butane torch + refill (€35), graphite crucible (€20). Total: ~€120

Reagents (for 40 servers): HCl 33% 2L (€8), H_2O_2 12% 2L (€6), $NaNO_3$ 500g (€12), $FeSO_4$ 500g (€8), sulfamic acid 500g (€10), NaOH granules 1kg (€6). Total: ~€50

Total first-batch investment: €245 (reusable for subsequent batches)

RATIO

◆ PROFITABILITY TABLE — BATCH OF 40 R640 SERVERS

Gross recoverable value: €880 to €2,300 (gold + copper + silver)

Reagent cost batch 1: €50

Equipment cost (amortised over 5 batches): €50/batch

Estimated processing time: 2 days (preparation + chemistry + neutralisation)

WEEE disposal cost avoided: €0 to €500 depending on provider

Net recovered batch 1: €780 to €2,200 — i.e. €19 to €55 per server

From batch 2: equipment is amortised. Marginal cost drops to €50 of reagents. Net recovered rises to €830 to €2,250.

◆ WHAT THE CALCULATION DOES NOT ACCOUNT FOR

The SysAdmin's time — 2 days of processing represents a salary cost. This process is relevant for an enthusiast who finds the process interesting in itself, or for an organisation processing sufficient volumes to justify the time investment. It is not a full-time business — it is opportunistic FinOps on hardware that would otherwise be a pure cost.

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · REGULATION — WHAT IS LEGALLY PERMITTED

LEGAL FRAMEWORK IN FRANCE

Recovering precious metals from IT hardware is legal in France in a private or professional context, subject to hazardous waste regulations. Three key points to know before starting.

◆ POINT 1 — EMPLOYER AGREEMENT FOR COMPANY EQUIPMENT

IT equipment belongs to the company until formally written off. Recovering servers off-inventory without written agreement creates legal risk. The correct approach: obtain a written free transfer deed or a signed decommissioning report from management or the procurement department. This document formalises the transfer of ownership and legally covers the recovery.

◆ POINT 2 — CHEMICAL WASTE PRODUCED

Acid solutions and metal sludge produced are classified as hazardous waste under European directive 2008/98/EC. For personal or lab use, volumes generally remain below mandatory declaration thresholds (10 litres of hazardous liquid waste per month for individuals). Above these thresholds, a declaration to the installations classées inspectorate (ICPE) may be required. Depositing sludge at a licensed specialist waste facility is always required — it is both a legal obligation and the right thing to do.

◆ POINT 3 — SELLING RECOVERED GOLD

Selling recovered gold is legal in France. Gold buyers (jewellers, buying counters, specialist smelters) buy industrial recovery gold. For quantities under a few grams, the transaction is informal. Above €500 in value per transaction, the buyer must verify the seller's identity. Keeping traceability of the source material (transfer deed) is good practice even below this threshold.

◆ NASSIHA — THIS PROCESS IS NOT DESIGNED FOR INDUSTRIAL SCALE

This process is documented for a SysAdmin enthusiast treating hardware from their own rack or their company's with agreement. For industrial volumes (tonnes of PCBs), licensed operators like Umicore, Aurubis or Comet Traitements have the infrastructure and authorisations — and far superior extraction yields. This study documents the DIY level — lab, torch, crucible, aquarium air pump.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

THE ENGINEER WHO MELTED THEIR GOLD UNDERSTANDS WHAT IS IN A SERVER

This process is more than a radical FinOps exercise. It is a lesson in applied chemistry, materials physics, and circular economy that the infrastructure engineer learns by doing. It completes the Foundation of Iron with a dimension that lectures cannot provide: the physical and chemical value of what they administer.

◆ WHAT THIS PROCESS TEACHES THAT A COURSE CANNOT

The engineer who has done this process once viscerally understands why PCIe connectors are gold and not copper — oxidation resistance, conductivity, ductility. They understand why DIMMs have "gold fingers" — perfect electrical contact over millions of insertion cycles. They know their Dell server contains gold in the same way an ingot does — differently concentrated, differently distributed, but real and measurable. This material knowledge is exactly what cloud training abstracts away and what the Foundation of Iron seeks to rebuild.

◆ THE LINK WITH "THE DIGITAL IRON"

"The Digital Iron" documented that the chip is the ingot of the 21st century — a rare, strategic, geopolitically sensitive resource. This process is the experimental demonstration of this thesis: the server the administrator installs literally contains precious metals, exactly as ore contains gold. The chain TSMC → chip → server → recovered gold is a real value chain that the DIY lab makes visible, tangible, measurable.

◆ FOR TRAINING AND CURRICULA

This process is a possible pedagogical format in BTS or DUT infrastructure curricula — a supervised practical lab day around dismantling, identifying high-value metal components, and supervised recovery. This is not a chemistry lesson — it is an infrastructure lesson that goes through chemistry to physically anchor the value of hardware. The student who has seen gold emerge from a PCIe connector understands its value differently from one who only read about it in a textbook.

RATIO

The FinOps of the last gram is not avarice. It is respect — for the physical value of what one administers, for the rare resources that have served, and for the engineer who knows what is in their rack down to the last atom.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

GRADE A REFURBISHED

What the second-hand machine says that the new one does not

RATIO

◆ CONTEXT OF THE STUDY

This study completes the Opération Dindon corpus on a hardware FinOps angle. It addresses Grade A refurbished server hardware as an infrastructure strategy — economically, ecologically and operationally. It draws on the author's direct experience as Head of SRE, on real purchase figures, and on a photograph of a machine that proves through its uptime counter that enterprise hardware is built to last.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION 1 · THE MACHINE WITH 4,122 DAYS OF UPTIME

THE BEST PROOF IS IN THE UPTIME COUNTER

4,122 days. That is the uptime counter of a legacy database machine, isolated in its corner of the rack, that nobody wanted to touch because it has been running for so long that everyone is afraid to stop it. 4,122 days is more than 11 years of uninterrupted operation.

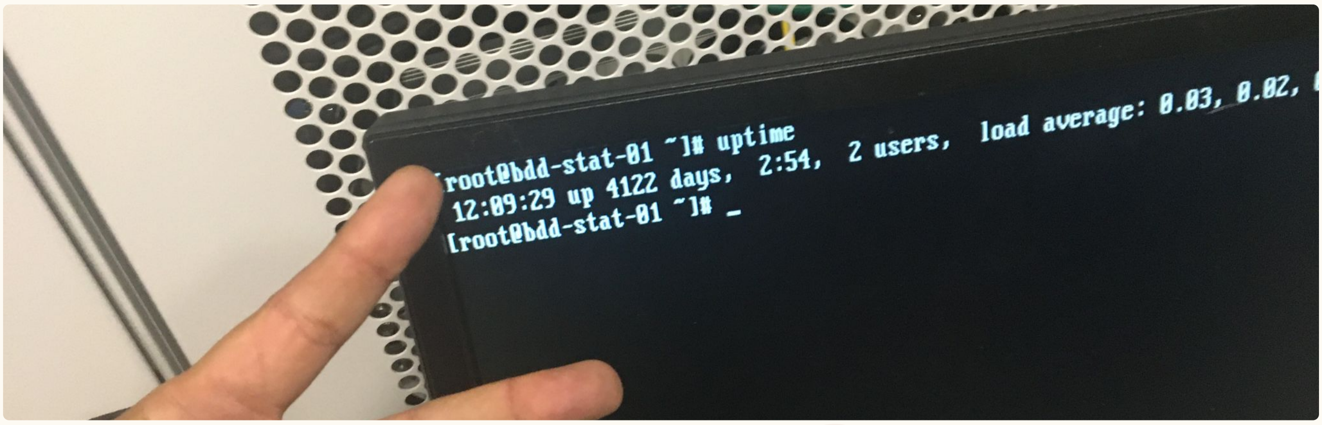

Uptime: 4,122 days — legacy database machine · © Amine RAITI

This machine is not an exception. It is the demonstration that enterprise server hardware is designed to last well beyond the 3-to-5-year renewal cycles that manufacturers recommend — and that finance departments approve without question. A server that has been running for 11 years without a major hardware incident is not a server on borrowed time. It is a server proving that the operational durability of enterprise hardware is structurally underexploited.

RAITI

◆ NASSIHA — WHAT THIS MACHINE DOES NOT DEMONSTRATE

4,122 days of uptime on an isolated legacy machine does not mean all hardware can be used indefinitely. This machine is a database under controlled load, without extreme demand, in a thermally stable environment. It demonstrates that the real operational lifespan of enterprise hardware is well above what standard procurement cycles assume — not that renewal is unnecessary.

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE NUMBER THAT DOES NOT LIE

€3,000 AGAINST €12,000 — IDENTICAL WORKLOAD

The comparison is simple. A Dell PowerEdge R640 Grade A refurbished with extended support — 192 GB RAM, 4×1 TB SSD, dual Xeon CPU — cost €2,100 in January 2024 from a specialist operator. In June 2026, with the surge in RAM component prices and general inflation, that same server trades at around €3,000. Its direct successor in new condition — a Dell PowerEdge R760 configured to equivalent specifications — is billed between €8,000 and €12,000 depending on configuration and procurement channel.

◆ THE 1-TO-4 RATIO

€3,000 against an average of €10,000: Grade A refurbished comes in at roughly one quarter of the new price for an identical workload. Across a fleet of 10 servers, the difference is €70,000. Across a fleet of 50 servers — a common size for a medium SRE infrastructure — it is €350,000. This is not a marginal saving. It is a budgetary architecture decision.

◆ WHAT THE PRICE INCLUDES ON THE REFURBISHED SIDE

The Grade A refurbished price is not the price of a bare machine. From serious operators it includes a full refurbishment (cleaning, testing of all components, replacement of defective parts), a contractual warranty, and the option of extended support covering hardware beyond manufacturer support. The comparison with new is therefore service for service — not a bare machine against a new one.

◆ THE PRICE OF THE XEON ALONE IS A REVELATION

On a new R750 (15th generation intermediate), a single Intel Xeon Gold 6346 processor is billed at $2,317 by the integrator. Two processors therefore represent over $4,600 before counting the chassis, RAM, storage and support. The R640 refurbished at €3,000 complete — dual CPU, 192 GB RAM, 4 SSDs, extended support — costs less than the two processors alone of an intermediate generation in new condition. This is not a promotion. It is the reality of the enterprise refurbished hardware market.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · WHAT GRADE A MEANS

NOT SECOND-HAND — REFURBISHED INFRASTRUCTURE

The word "second-hand" evokes risks that Grade A refurbished in an enterprise environment does not carry. The distinction is structural and deserves to be made explicit — because it is often the confusion between the two that blocks purchasing decisions in organisations that have never tested the model.

◆ WHAT GRADE A MEANS CONCRETELY

Full refurbishment: the machine is fully disassembled, cleaned, all components individually tested, defective parts replaced with certified original parts. It leaves in an internal condition equivalent to new.

Grade A cosmetic: the machine may show minor signs of use on the chassis — light scratches, coating wear. It does not look like a machine fresh from the box. In a rack, behind a bay door, this has zero operational relevance.

Contractual warranty: serious operators offer a 12-to-36-month warranty on refurbished hardware, with replacement of defective components.

Extended support: specialists such as Evernex, Park Place Technologies, Curvature or MaintenX offer multi-vendor maintenance contracts covering Dell, HP, Cisco, NetApp and others — with SLA, on-site intervention and parts replacement — on hardware beyond official manufacturer support. This extended support is the keystone that makes refurbished viable in critical production.

◆ WHAT DISTINGUISHES A GOOD OPERATOR FROM A BAD ONE

The quality of Grade A refurbished depends entirely on the operator. A good operator publishes their testing procedures, provides the Dell service tag (verifiable on Dell's support site for the machine's full history), offers a clear support SLA, and holds spare parts stock for the models they sell. A bad operator sells cleaned hardware without functional testing under a superficial warranty. Due diligence on the operator is the real step in securing a refurbished purchase — not systematic distrust of the model.

◆ NASSIHA — REFURBISHED IS NOT FOR EVERYONE

Grade A refurbished with extended support is suited to stable production workloads on proven hardware. It is not suited to workloads requiring the latest processor generations (AI, HPC), architectures dependent on PCIe Gen5 or DDR5 features, or environments where direct manufacturer support is a contractual or regulatory requirement. For the rest — and that is the majority of infrastructure workloads — it is the right choice.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE ECOLOGICAL ARGUMENT

EVERY MACHINE NOT BOUGHT NEW IS A CHIP NOT ORDERED FROM TAIWAN

The ecological argument for refurbished is often presented as secondary — a good conscience on top of a good price. This is a framing error. The ecological argument is an industrial sovereignty argument as much as an environmental one, and it connects directly with the theses of "The Digital Iron" in the Opération Dindon corpus.

◆ MANUFACTURING A SERVER: WHAT IT CONSUMES

Manufacturing a new server is one of the most resource-intensive acts in the digital industry. It mobilises rare earths extracted in China, chips manufactured in Taiwan at TSMC, printed circuit boards produced in Southeast Asia, and an assembly process consuming significant energy and water. The carbon footprint of server manufacturing represents a substantial proportion of total lifetime carbon footprint — some studies estimate it between 50 and 80% of the total over 4 years.

◆ EXTENDING LIFESPAN: THE MOST EFFECTIVE LEVER

A server used for 8 years instead of 4 halves its manufacturing energy consumption per year of service. A refurbished server used for a second life of 4 years avoids entirely the manufacture of a new server. Across a fleet of 50 servers renewed every 4 years, switching to Grade A refurbished represents avoiding the manufacture of 50 new servers every 4 years — with everything that implies in terms of TSMC chips not ordered, rare earths not extracted, and manufacturing cycles not triggered.

◆ THE LINK WITH "THE DIGITAL IRON"

"The Digital Iron" documented that 90% of the world's advanced chips are manufactured by TSMC in Taiwan, on a territory of 35,000 km² exposed to structural geopolitical tension. Every new server not purchased is an order of chips not placed with TSMC. Grade A refurbished is therefore also a lever for reducing dependence on the semiconductor supply chain — not only a cost-reduction lever. These two arguments arrive at the same conclusion by different paths.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHAT MAKES REFURBISHED VIABLE IN CRITICAL PRODUCTION

The main barrier to adopting refurbished in critical production is not hardware reliability — that is demonstrated. It is the support question: what happens when a component fails on hardware beyond manufacturer support? The answer is third-party extended support — a structured market, with established actors, contractual SLAs and spare parts stock that allow hardware to be maintained long after the end of official support.

◆ WHAT AN EXTENDED SUPPORT OPERATOR OFFERS

Specialist operators — among them Evernex, Park Place Technologies, Curvature and MaintenX — offer multi-vendor maintenance contracts covering Dell, HP, Cisco, NetApp and others. These contracts include an intervention SLA (4h, 8h, NBD depending on the level chosen), certified spare parts stock, 24/7 on-call support, and technical assistance. The cost of such a contract is significantly below Dell's ProSupport manufacturer support on new hardware — while covering hardware that Dell no longer officially supports.

◆ SERVICE TAG VERIFICATION — A MANDATORY STEP

Before any refurbished purchase, the machine's service tag must be verified on Dell's support website. This unique number reveals the original manufacturing date, the history of warranty interventions, and any parts replaced. A transparent service tag is the first sign of a serious operator. An operator who refuses to provide it before purchase is an operator to avoid.

◆ THE 3-YEAR ROI — THE CALCULATION THAT CONVINCES CIOs

Over 3 years, for 10 Grade A refurbished R640 servers at €3,000 each with extended support at €500/year/machine:
Total investment: €30,000 (hardware) + €15,000 (3-year support) = **€45,000**
Over 3 years, for 10 new R760 servers at €10,000 each with Dell Basic ProSupport:
Total investment: €100,000 (hardware) + €18,000 (support) = **€118,000**
Saving: **€73,000 over 10 machines over 3 years** — for an identical workload.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

GRADE A REFURBISHED FIRST — A PROCUREMENT POLICY

The proposal of this study is not to replace all new hardware with refurbished. It is to reverse the burden of proof in purchasing decisions: instead of justifying why you choose refurbished, justify why you choose new. New hardware has legitimate use cases. They are fewer than commonly assumed.

◆ WHEN TO CHOOSE GRADE A REFURBISHED

— Stable production workloads on proven technologies: databases, web servers, virtualisation, storage
— Environments where direct manufacturer support is not a contractual or regulatory requirement
— Fleet renewal on well-documented server generations (R630, R640, R740, HP DL360 Gen9/Gen10)
— Lab, staging, development environments — without exception

◆ WHEN TO CHOOSE NEW

— Workloads requiring the latest processor generations: AI, HPC, GPU inference
— Architectures dependent on PCIe Gen5, DDR5, or features unavailable on refurbished generations
— Contractual or regulatory requirements mandating direct manufacturer support
— Edge computing and space-constrained environments where latest-generation density is determinative

◆ THE THREE CONCRETE MEASURES

Measure 1: integrate a refurbished vs new analysis into every hardware procurement process exceeding €10,000. The decision to choose new must be technically justified, not assumed by default.
Measure 2: integrate real operational lifespan and carbon footprint into hardware procurement evaluation criteria — on a par with performance and acquisition cost.
Measure 3: qualify at least one third-party extended support operator in the organisation's supplier panel. This pre-qualification removes the support barrier and makes refurbished immediately actionable when the decision is made.

RATIO

4,122 days of uptime. The machine does not lie. Enterprise hardware lasts far longer than procurement cycles assume. Grade A refurbished is the economic, ecological and operational proof.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE STRATEGIC REFURBISHED

11× Cheaper Than Cloud · And Geopolitical Insurance

Why Owning Your Hardware Is a Strategic Decision

RATIO

◆ THE THESIS

Refurbished Grade A was documented in "The €3,000 VM and the €2,800 SRE" as a FinOps choice — 11 times cheaper than the GCP VM over 5 years. This study adds the dimension the previous one did not treat: Refurbished is not just a FinOps choice. It is geopolitical insurance whose value increases exactly when cloud becomes most vulnerable. Both properties together — economic and geopolitical — make Refurbished Grade A the most defensible decision in the corpus.

FINOPS RATIO	INDEPENDENCE	WATERMARK
11×	**TSMC INDEPENDENT**	**RATIO**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION 1 · THE TWO PROPERTIES OF REFURBISHED GRADE A

FINOPS: 11× CHEAPER · GEOPOLITICAL: INDEPENDENT OF FUTURE PRODUCTION

◆ PROPERTY 1 — FINOPS (RECAP FROM PREVIOUS STUDY)

A Refurbished Grade A Dell PowerEdge R750 (32 cores, 512 GB RAM, 10 TB NVMe) costs €4,000 to €6,000. Depreciated over 3 years with opex (power, network, colocation): approximately €270/month. Over 3 years: €9,720. The equivalent GCP VM (n2-highmem-32) over the same period: €108,000. Ratio: 11×. This property was documented in "The €3,000 VM and the €2,800 SRE".

◆ PROPERTY 2 — GEOPOLITICAL (NEW IN THIS STUDY)

A Refurbished Grade A server bought today contains chips fabricated 2 to 3 years ago. These chips already exist in the physical world — they have been produced, tested, assembled. Their availability no longer depends on:
— Future TSMC production in Taiwan
— US export restrictions on advanced chips
— ASML EUV machine availability
— The geopolitical situation in the Taiwan Strait
— Hyperscaler procurement policy for new GPU servers
These chips are there. They work. They will continue to work for 5 to 7 additional years with standard maintenance. Their physical existence is established. It cannot be revoked by a future geopolitical event.

◆ THE FORMULA LINKING THE TWO PROPERTIES

Refurbished Grade A is 11 times cheaper than the GCP VM in a stable world. And in an unstable world — Taiwan crisis, export restrictions, cloud capacity rationing — its relative value increases further. It is the asset whose strategic value grows exactly in the scenarios where the cloud alternative is most exposed.
These two properties are not independent. They are two faces of the same fundamental decision: own your hardware or rent someone else's.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · WHAT HAPPENS IN A CRISIS — COMPARATIVE SCENARIO

100% CLOUD · BARE-METAL · HYBRID — IN CRISIS

100% CLOUD ORGANISATION — IN A TAIWAN CRISIS

Infrastructure runs on hyperscaler servers. In a crisis, the hyperscaler cannot order new GPU servers — TSMC is down. It maintains existing capacity but cannot extend it. If the organisation needs to scale, it hits rationing. It is in the queue — after the hyperscaler's own needs, after Enterprise Agreement customers. Without guaranteed contractual priority, it waits.
It also cannot repatriate its infrastructure — it no longer has the skills (documented in The Newspeak), the engineers (−34% Linux jobs), and migration takes 6 to 12 months (documented in The Accompanied Exodus).

REFURBISHED BARE-METAL ORGANISATION — IN A TAIWAN CRISIS

Infrastructure runs on physical servers it owns. These servers were purchased and delivered before the crisis. Their chips are there, functional, independent of future production. In a crisis, it can maintain normal activity. It can even purchase additional Refurbished on the secondary market — whose value rises with chip scarcity but remains available since the world's existing server park is massive. It is outside the cloud rationing queue.

HYBRID ORGANISATION (BARE-METAL + CLOUD) — IN A TAIWAN CRISIS

Core infrastructure runs on its own servers — critical services, sensitive data, baseline workloads. Cloud is used for temporary peaks or non-critical services. In a crisis, it can rebalance cloud workloads to physical infrastructure. It is not entirely dependent on the vendor's queue. This is the model documented in "The Sovereign Interface" — the Anti-Corruption Layer enabling switching without application rewrite.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · SYNTHESIS — THE DECISION WITH TWO JUSTIFICATIONS

SAME DECISION · TWO REASONS · ONE CONCLUSION

DIMENSION	100% CLOUD (ON-DEMAND VM)	REFURBISHED GRADE A
5-year cost	**€180,000 (residual value: €0)**	**€16,200 (residual value: €500–1,500)**
FinOps ratio	Reference (1×)	**11× cheaper**
TSMC / future production dependency	Total (hyperscaler depends on TSMC)	None (chips already produced)
Continuity in silicon supply crisis	Rationing by contractual priority	Independent — outside the queue
Relative value in supply crisis	Decreases (constrained availability)	Increases (new chip scarcity)
Optimal decision in both scenarios	Stable world: correct but expensive. Unstable world: exposed.	**Stable world: 11× cheaper. Unstable world: geopolitical insurance.**

Refurbished Grade A is the only infrastructure decision
that is optimal in both scenarios — stable and unstable.
That is why "owning your hardware" is not nostalgic sovereigntism.
It is elementary risk management.
And it is 11 times cheaper too.

NEMO SUPRA LEGEM EST

RAITI

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE GENTLE BCP

Business Continuity as a Muscular Reflex

Why the Hyperscaler BCP is an Architecture

and the Managed Service BCP is a Reflex

RATIO

◆ **THE THESIS**

The Business Continuity Plan is the CIO's personal responsibility when infrastructure goes down. Hyperscalers sell it as an automatic architectural property. Sovereign managed service providers practise it as a repeated muscular reflex. This study documents the difference between these two realities — and why it is decisive for organisations that want real business continuity, not documented business continuity.

FORMULA	ANGLE	WATERMARK
ARCHI VS REFLEX	**RATIO COMPARATIVE**	**RATIO**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · THE COMFORT LIE — WHAT THE HYPERSCALER BCP REALLY SAYS

MULTI-AZ · MULTI-REGION · AUTOMATIC FAILOVER — AND YET

Hyperscalers sell BCP as a native architectural property. Multi-AZ means your resources are spread across multiple availability zones in the same region. Multi-region means your data is replicated to another geographic region. Automatic failover means the switchover happens without human intervention. These three promises are true — under certain conditions, for certain services, according to certain configurations. What they do not say is more important than what they say.

◆ WHAT MULTI-AZ DOES NOT SAY

Availability zones share the same regional network backbone. In December 2021, an AWS us-east-1 incident took down theoretically "multi-AZ" services because the control plane (the service that orchestrates failover) was itself centralised and affected. Multi-AZ protects against a single physical datacentre failure. It does not protect against control plane failures, regional network failures, or shared management service failures.

Furthermore, multi-AZ costs. Every duplicated resource is billed. Traffic between AZs is billed. An RDS Multi-AZ instance costs approximately 2x a Single-AZ instance. A complete multi-AZ architecture can cost 1.5x to 2x the primary infrastructure — not counting replication egress fees.

◆ WHAT MULTI-REGION DOES NOT SAY

Multi-region replication is real but complex. It requires solving the data consistency problem: in synchronous mode (zero data loss), inter-region latency degrades performance. In asynchronous mode (non-zero RPO), one must accept a data loss window. This architectural choice is not trivial — it depends on the nature of the data, the business, and regulatory requirements.

Implementation requires mastery of Route 53 (DNS failover with health checks), Aurora Global Database or DynamoDB Global Tables, S3 Cross-Region Replication, CloudFront, WAF, ACM (certificates in each region), and dozens of other services — each with its own subtleties, failure modes and hidden costs. A well-built multi-region AWS BCP costs between 2x and 5x the primary infrastructure cost.

RAITI

◆ THE FAILOVER TEST NOBODY DOES

Testing a multi-region failover in production is high-stakes surgery. It requires a maintenance window, team coordination, a rollback plan. It is expensive — test traffic is billed like production traffic. And it is frightening — a configuration error can worsen the incident instead of resolving it.

Result: in most organisations, the multi-region BCP is only tested in staging — which is never identical to production. When the real incident arrives at 3am, DevOps discover that the failover site's Terraform has not been updated since the last database migration. That the secondary region's load balancer SSL certificate expired three weeks ago. That the DNS TTL is still set to 3600 seconds.

Amine RAITI · CC BY-NC-SA 4.0

THE PHYSICAL FAILOVER · THE LIVING RUNBOOK · ADAPTATION TO REAL NEEDS

An experienced managed service provider has two datacentres — often connected by a dedicated fibre they own or lease at the physical level, outside the public Internet backbone. They know the power supplies, the UPS units, the generators, the core switches, the border routers. They built this infrastructure. They know what happens at every layer when a failover is triggered.

◆ THE MASTERED PHYSICAL CHAIN

The managed service failover begins at the physical level — where the hyperscaler BCP begins at the software level. When datacentre A loses its main power supply, the UPS takes over (milliseconds), then the generator starts (30 to 60 seconds). During this time, network flows are rerouted to datacentre B via the dedicated fibre — not via the Internet backbone, whose availability is outside the provider's control.

This physical layer mastery is what cloud cannot offer. When the problem is a cut transoceanic cable, a major BGP incident, or a saturated regional backbone — the cloud's software BCP does nothing. The managed service provider who controls their dedicated fibre has resilience orthogonal to Internet infrastructure incidents.

◆ THE BCP ADAPTED TO REAL APPLICATION NEEDS

The managed service provider does not propose a BCP services catalogue — they design a continuity architecture adapted to the client's real needs. This conversation starts with fundamental questions:

RPO (Recovery Point Objective): how much data can be lost? Billing database: RPO zero (synchronous, no loss). Application logs: RPO 4h acceptable (asynchronous). User sessions: RPO 15 minutes (semi-synchronous).

RTO (Recovery Time Objective): how quickly must operations resume? Website: RTO 5 minutes (DNS failover + HAProxy). Back-office: RTO 30 minutes. Batch processing: RTO 4h.

These nuances are implemented in custom scripts, bespoke HAProxy configurations, MySQL or PostgreSQL replication rules specific to each service's RPO. Not a checkbox in a catalogue — a tailored architecture.

RATIO

◆ THE LIVING RUNBOOK — NOT THE DOCUMENT NOBODY READS

The managed service team has a failover Runbook for every incident type. Not a bureaucratic document — an operational guide tested and updated after every real or simulated incident. It says exactly: who does what, in what order, with what command, with what success criterion. And it notes the traps: service X always takes 3 minutes longer than expected, service Y requires a manual restart if the failover lasted more than 10 minutes.

Amine RAITI · CC BY-NC-SA 4.0

AN ANNUAL COPIL WITH REAL FAILOVER · NOT A POWERPOINT ON RESILIENCE

◆ THE HYPERSCALER TEST — WHAT ACTUALLY HAPPENS

In most cloud organisations, BCP testing follows this scenario:

Month 1: the team decides to test failover. They realise staging is not up to date.

Month 2: staging updated. Test in staging. Partially works — 3 of 12 services do not failover correctly.

Month 3: fixes. Re-test in staging. Better, not perfect.

Month 4: decision not to test in production this year — "too risky before the holidays".

Month 12: real incident. Production Terraform has diverged from staging for 4 months. Secondary database has 6 hours of replication lag. Secondary region SSL certificate expired 3 weeks ago. Real RTO: 4 hours instead of the promised 15 minutes.

◆ THE MANAGED SERVICE TEST — THE FAILOVER COPIL

A reputable managed service provider integrates BCP testing into the client's contractual lifecycle. Not as a one-off project — as a recurring, planned event with a dedicated Steering Committee.

Typical failover test schedule:

D-30: client notified of date and maintenance window. Verification that all target services are documented and the Runbook is current.

D-7: technical review — which services failover in what order, what are the success criteria, what is the rollback procedure.

Day of test, maintenance window: real production failover. The client is present (in person or remote). They observe. They validate. They measure the real RTO.

After failover: post-incident COPIL. Which services took longer than expected? Why? What changes in the Runbook? The Runbook is updated before end of week.

The client goes home with a conviction: their BCP works. Not with an architecture diagram.

RATIO

◆ THE DEFINITIVE FORMULA

The hyperscaler BCP is an architecture documented in a Confluence nobody reads, tested in staging that is never identical to production, validated by a checklist nobody has filled in for 8 months.

The managed service BCP is a muscular reflex repeated until the team can do it with their eyes closed, documented in a living Runbook updated after every incident, validated by a COPIL where the client watched their production failover and return.

The hyperscaler BCP is a promise. The managed service BCP is a demonstration.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · COMPARISON TABLE — HYPERSCALER BCP VS MANAGED SERVICE BCP

TERM BY TERM · DIMENSION BY DIMENSION

DIMENSION	HYPERSCALER BCP	MANAGED SERVICE BCP
Nature	Software architecture (proprietary managed services)	Muscular reflex (mastered physical procedure)
Physical layer	Not mastered — depends on AWS/GCP/Azure backbone	Mastered — dedicated fibre, power supply, cooling
Cost	2x to 5x primary infrastructure + replication egress fees	Included in contract or fixed predictable fee
RPO/RTO	Defined by managed service capabilities	Defined by real client business needs
Failover test	Rare, costly, frightening. Often done in staging only.	Planned, contractual, done in production with client.
Documentation	Confluence not updated since last migration	Living Runbook, updated after every incident or test
Application adaptation	Checkbox in a catalogue (Aurora Global, DynamoDB Global)	Custom scripts, HAProxy, MySQL/PG replication per RPO
Network cable incident	Software BCP useless if the problem is physical (backbone)	Dedicated fibre outside backbone — orthogonal resilience
Result for the CIO	A promise in an architecture diagram	A conviction gained from having watched their production failover

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE GENTLE BCP — WHY "GENTLE"?

NO RUPTURE · NO PANIC · NO SURPRISE

The word "gentle" in "Gentle BCP" is not a metaphor — it is a description of the nature of a well-designed managed service failover. Gentle as opposed to brutal. The hyperscaler failover during a real incident is brutal — services go down one after another in an unpredictable order, teams chase logs, decisions are made under pressure with incomplete information. The well-prepared managed service failover is gentle — it follows a predefined order, it is executed by people who have done it before, and it has a documented rollback if something does not go as planned.

◆ THE THREE PROPERTIES OF THE GENTLE BCP

Progressive: the failover happens service by service, in dependency order. First the database (source of truth), then the APIs that consume it, then the frontend. Each step is validated before moving to the next. There is no "everything failovers simultaneously" — because simultaneous failover creates unpredictable timing problems.

Reversible: each failover step has a documented rollback procedure. If service Y does not start correctly on site B, the team knows exactly how to return to site A without worsening the situation. This reversibility is tested — not assumed.

Observed: the failover is observed in real time by the provider's and client's teams. Key metrics are monitored (latency, error rate, active connections). "Continue" or "rollback" decisions are made on real data, not assumptions.

◆ THE CORPUS LINK — THE ACCOMPANIED EXODUS AND THE GENTLE EXIT

The Gentle BCP is the business continuity declension of two corpus theses. "The Gentle Exit" proposes progressive migration without rupture — the Gentle BCP applies the same principle to emergency failover. "The Accompanied Exodus" proposes a managed service provider as migration guide — the Gentle BCP makes the same provider the guide of continuity.

The organisation that migrated to the managed service provider via the Accompanied Exodus has, as a bonus, a Gentle BCP — because the provider knows the infrastructure they themselves migrated, wrote the TAD, and integrated the BCP into the initial design rather than as a later addition.

RATIO

◆◆◆

The CIO who watched their production failover and return
during a contractual test with their managed service provider
sleeps better at night than the one who read
their AWS multi-region architecture diagram.
One has a promise. The other has a demonstration.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

PUBLIC DOCUMENT · OPERATION DINDON · JUNE 2026 · PLAIN-LANGUAGE

THE INVISIBLE PLUMBING

What nobody shows you behind your screen

HUMAN

◆ THE CENTRAL THESIS

We say "it's in the cloud." The image is lovely. It is also completely false. Nothing you do online floats in the air: a hard drive spins, a cable carries a signal, an air-conditioned building processes your request. This document shows you that invisible plumbing — no complicated figures, no international law, just examples from your everyday life.

PAGES	EXAMPLES	WATERMARK	OWNERS
8	4	HUMAN DOMINANT	3

Amine Raiti — Infrastructure Architect & SRE

Former engineering school professor · Trainer since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

BEHIND EVERY DIGITAL GESTURE, A PHYSICAL MACHINE SWITCHES ON

We say "it's in the cloud." The image is lovely. It is also completely false. Nothing you do online floats in the air. When you tap your screen, somewhere on Earth, a physical machine switches on. A hard drive spins. A cable carries an electrical signal. An air-conditioned building, full of servers running hot day and night, processes your request.

◆ THE IMAGE TO REMEMBER: TAP WATER

You turn on the tap, water flows. You never think about the water towers, the pipes buried under the city, the pumping stations. Yet without them, not a single drop would arrive. Digital services work exactly the same way: a smooth gesture on screen — and behind it, a giant industrial plumbing system nobody ever sees.

This document shows you that plumbing. Not with complicated figures, not with international law. Just with examples from your everyday life.

◆ A NOTE ON SCOPE

This document draws on a French research corpus (Opération Dindon), so the examples — Doctolib, the CAF, French datacentre locations — are French. But the mechanism described here is not a French problem. The same three companies operate the same kind of factories near Frankfurt, Dublin, Singapore, São Paulo, Mumbai, and dozens of other cities. Wherever you live, replace "Doctolib" with your own country's medical booking platform, "the CAF" with your own welfare agency, and the reasoning holds exactly the same.

◆ WHAT TO REMEMBER

"Digital technology has nothing of a magical cloud about it: every time you touch your screen, a physical machine switches on, heats up, and consumes electricity somewhere on Earth."

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

I

EVERYDAY EXAMPLE · ENTERTAINMENT

TONIGHT, A SERIES ON NETFLIX

◆ THE GESTURE — 0.3 SECONDS

You press "Play". The picture appears, smooth, instant. No effort. No waiting.

◆ THE PHYSICAL REALITY — A FEW GIGABYTES

A single episode is several gigabytes of data. That weight doesn't teleport: it travels, at very high speed, through fibre-optic cables buried under roads, from a hard drive sitting in a specific building — all the way to your living room. This is not a file that "exists somewhere in the ether." It is a file with a real weight, stored on a real disk, that must physically travel a real distance to reach you.

◆ THE REMOTE CONTROL, NOT THE MEMORY

Your phone or your TV is just a screen. All the intelligence — the video file, the storage, the computation — sits elsewhere, in a digital factory. Your device is a very sophisticated remote control. Nothing more.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

2

EVERYDAY EXAMPLE · HEALTH

A DOCTOR'S APPOINTMENT

◆ THE GESTURE — 2 CLICKS

You click an open, green time slot, Tuesday at 2:30pm. A confirmation appears. The appointment is booked.

◆ THE PHYSICAL REALITY — ONE LINE, ONE DISK

That click becomes a line of text written onto a physical hard drive, in a specific building. The doctor will check that very same information on their own screen on Tuesday morning — to know who to see, when, and why. If that hard drive fails, if the building housing it loses power, that line of text no longer exists anywhere else. The appointment, literally, vanishes.

◆ REAL ADDRESSES, NOT ABSTRACT PLACES

The big American cloud companies (Amazon, Microsoft, Google) have indeed built digital factories in France. They have precise addresses:

La Courneuve · Vitry-sur-Seine · Clichy · Paris-Saclay · Wissous · Marseille · Ferrières-en-Brie

The building is French. But the rules, the prices, and the access keys are still decided elsewhere. The same logic applies everywhere — these companies have built identical factories near Frankfurt, Dublin, Singapore, São Paulo, and dozens of other cities worldwide. France is simply the example used here; replace these names with the ones near you, and the same reasoning holds.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

3

EVERYDAY EXAMPLE · CITIZENSHIP

A WELFARE OR HEALTH-FUND PROCEDURE

◆ THE GESTURE — ONE FORM

You fill in a form and click "Submit". The site confirms: your request is registered.

◆ THE PHYSICAL REALITY — A MACHINE, A BUILDING

That click sends an electrical impulse to a machine — owned by the State or a service provider — which records your request on physical storage. Your rights, your file, your benefit: all of it exists somewhere, as data stored on real hardware. Even your administrative relationship with the State — a fundamental right — runs through these same factories of concrete and cable.

◆ CONCRETE, NOT DUST

We sometimes picture computing as something fragile — a small, dusty desktop computer. The reality is the opposite: these are massive buildings, air-conditioned around the clock, protected by backup power generators, almost comparable to heavy industrial infrastructure.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

4

EVERYDAY EXAMPLE · MONEY

A €20 TRANSFER TO A FRIEND

◆ THE GESTURE — 3 SECONDS

You open the app, type "20", confirm. The money seems to "leave" instantly.

◆ THE PHYSICAL REALITY — A REMOTE CONTROL, A DISTANT COMPUTER

Your money has not, for a long time, been sitting in brick vaults. It is a number, written and edited on the hard drives of giant computers, somewhere. Your smartphone app is just a visual remote control: it sends an instruction to a physical machine, which actually carries out the operation.

◆ THE QUESTION THAT MATTERS

"If a single company owns every pumping station in the city, it can decide tomorrow to double the price of your water bill or change the rules of access, and you will have nowhere else to go. That is what losing everyday sovereignty means: depending on the goodwill of a single owner for the most ordinary gestures."

OPERA

Amine RAITI · CC BY-NC-SA 4.0

A SIMPLE QUESTION

IF YOU DON'T KEEP YOUR OWN PLUMBING, SOMEONE ELSE TAKES IT

These buildings, these cables, these hard drives do not belong to nobody by chance. Someone built them, someone owns them, someone sets the rules. And a simple observation follows: a technical space left vacant never stays that way. If a country or a company doesn't build, maintain, and train enough people to run its own digital plumbing, someone else builds it instead — and keeps the keys.

◆ IT'S NOT JUST ANYONE WAITING

Today, this global plumbing isn't shared out among dozens of different players. It is mostly owned by three American companies — Amazon (AWS), Microsoft (Azure) and Google — who already own most of the world's giant digital factories. The more a country or a company gives up control of its own infrastructure, the smaller the share that remains free — and the more likely it is that whatever gets built next falls, by default, into those same three hands.

◆ A DETAIL THAT SAYS A LOT

While we're told servers and cables are "old world" technology, those same three companies spend tens of billions of dollars every single year building... servers and cables. Concrete, copper, fibre, land, dedicated power plants: they are pouring money into exactly what others have stopped funding. If it were a dying sector, it wouldn't be worth that much. They know exactly what they're doing ;)

◆ WHY THREE OWNERS IS DIFFERENT FROM THIRTY

If a hundred different companies each owned a small piece of your city's water network, losing one of them wouldn't change much: ninety-nine others would remain. But if only three already own nearly the entire network, their decisions — a price that changes, a rule that tightens, an outage at one of them — instantly affect a huge share of everything that depends on water. That is exactly what is at stake with digital infrastructure: the fewer different owners there are of the plumbing, the more each one matters — and the more serious it becomes to also hand them whatever is still left to build.

LEX

Amine RAITI · CC BY-NC-SA 4.0

IN CLOSING

LOOKING AT DIGITAL THROUGH ITS MACHINES

Netflix, Doctolib, the welfare office, your bank. Four ordinary gestures. Four times, the same reality: behind the screen, hardware, cables, buildings, somewhere — owned by someone.

◆ WHAT THE IRON NEVER LIES ABOUT

The Iron does not lie. A hard drive spins or it doesn't. A cable carries a signal or it doesn't. This stubborn, silent material reality is the invisible foundation of every digital gesture you have made since this morning.

You cannot change Internet's plumbing yourself.

But the next time you touch your screen to watch a show or book an appointment,

you will know that behind the visual magic, a physical factory has just kicked into gear.

Looking at digital through its machines is the beginning of understanding the real world.

Amine RAITI · Operation Dindon · 2026

NEMO SUPRA LEGEM EST

CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

أَتَسْتَبْدِلُونَ
الَّذِي هُوَ أَدْنَى
بِالَّذِي هُوَ خَيْرٌ

Are You Trading What Is Better For What Is Inferior?

Sarcastic Reflections on the Cloud Paradox

Cold aisle — server racks · Property of Amine RAITI

◆ THE QUESTION

You have this. AWS has this. The difference: you know yours. You know what is inside, who cabled them, how they work. Their racks — you will never see them, you will not know exactly where they are, and if someone knocks on the door with a warrant, you will be the last to know. Why are you trading what is better for what is inferior?

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · THE FOUNDING PARADOX — THEY MASTER WHAT YOU ARE ABANDONING

AWS HAS 70% OF THE CLOUD MARKET · THEIR ADVANTAGE: BARE-METAL · YOUR STRATEGY: ABANDON BARE-METAL

AWS, GCP and Azure control roughly 70% of the global cloud market. Their competitive advantage — the one that distinguishes them from every competitor — is absolute mastery of bare-metal. They know how to cable a datacentre, size a power supply, choose a processor, optimise cooling, build a physical network at planetary scale. They have thousands of engineers doing exactly what the corpus calls "the mastery of iron". It is on this mastered physical substrate that they build the cloud services they sell.

◆ THE MOST ELEGANT BUSINESS MODEL IN THE HISTORY OF COMPUTING

Step 1: master bare-metal perfectly.
Step 2: convince your customers that bare-metal is "legacy", "complex", "not your core business".
Step 3: sell them access to your bare-metal as a managed service, with a comfortable margin, a "noncancellable" contract, and California jurisdiction.
Step 4: watch your customers let go of their bare-metal engineers — eliminating any possibility of going back.
Brilliant. Truly. The corpus calls this "The Anatomy of the Loss". Hyperscalers call it "the cloud".

◆ TRANSLATED INTO PLAIN LANGUAGE

You have a cold aisle. Racks. Engineers who know what is inside. Someone convinces you that this is a problem. You decide to hand everything over to people who have exactly the same thing — but much larger, with much less transparency, under US law, with a margin that would make any Paris estate agent blush. Then you pay egress fees to retrieve your own data.
أَتَسْتَبْدِلُونَ الَّذِي هُوَ أَدْنَىٰ بِالَّذِي هُوَ خَيْرٌ
Are you trading what is better for what is inferior?

◆ NASSIHA — THIS IS NOT A CRITIQUE OF CLOUD IN GENERAL

Cloud has real advantages — elasticity, no upfront investment, complex managed services. The question is not "cloud or no cloud". The question is: why go to extraterritorial actors who make you dependent, when local and national actors offer the same services, under European law, with margins that allow real local R&D?

Amine RAITI · CC BY-NC-SA 4.0

YOU WERE THE OWNER · YOU BECOME THE TENANT · YOU PAY THE SUB-LANDLORD'S MARGIN

Imagine the following scenario. You own a flat. It works. You know it. Your plumber knows it. Your electrician knows it. Then someone convinces you that managing a flat is complicated, that it is not your core business, that you should hand it over to a professional. You give them the keys. They sub-let your own flat back to you at three times the price, with a three-year non-cancellable lease, fees if you want to get your belongings back, and a clause saying that if the US justice system wants to enter, it can do so without telling you.

◆ THE MECHANICS IN NUMBERS — WHAT THE MARGINS MEAN

A good quality physical server — 2 sockets, 512 GB RAM, 10 TB NVMe — costs approximately €15,000 to €25,000 to buy. Amortised over 5 years with colocation and power: approximately €500 to €800/month.

The cloud equivalent at AWS (r6i.16xlarge instance, 64 vCPU, 512 GB RAM): approximately €3,500 to €4,500/month on-demand. As a 3-year Reserved Instance: approximately €1,800 to €2,200/month.

Ratio: 2.5x to 4x the bare-metal cost. Over 5 years, for a single server, the difference is €80,000 to €200,000. This difference is the hyperscaler's margin — and egress fees are not included.

◆ WHAT THIS MARGIN FINANCES — FOR THEM

With European companies' cloud budgets, AWS finances new datacentres, new regions, new proprietary services that create new dependencies. GCP finances Gemini and its AI models. Azure finances Microsoft ecosystem integration.

With a tiny fraction of these same budgets — redirected to local providers — Ecritel, DRI, OVH, Scaleway, Infomaniak could fund R&D that would make Europe less dependent. They could build equivalents of SQS, Pub/Sub, DynamoDB — under European law, with European engineers, in datacentres whose cold aisles you could one day visit.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

ECRITEL · DRI · OVH · SCALEWAY · INFOMANIAK · IONOS · CLARANET — THEY EXIST

In France and across Europe, there are very good quality cloud providers, subject to European law, with local technical teams, capable of building custom solutions, understanding sector-specific regulatory constraints (health, finance, defence), and whose margins would remain in Europe. They are called Ecritel, DRI, OVH, Scaleway, Claranet, Infomaniak, IONOS — and a dozen others. They have cold aisles. They have racks. They have engineers.

◆ WHAT THEY CAN DO THAT THE BROCHURES DO NOT SAY

Custom work: a local provider can adapt their infrastructure to your specific needs. A hyperscaler adapts you to their infrastructure. The difference is not cosmetic — it is architectural.

Equivalent services: equivalents of SQS, Pub/Sub, DynamoDB, S3 exist at European providers — or can be built on open source technologies (Kafka, MinIO, ScyllaDB, RabbitMQ) hosted there. The Sovereign Interface documented in the corpus allows switching without application rewrite.

Native regulatory compliance: a French datacentre hosting health data is subject to HDS. A European datacentre hosting financial data is subject to DORA. No need to build a compliance layer on top of a service designed for the US market.

The possible visit: you can ask to visit their cold aisle. Try that with AWS.

◆ THE R&D ARGUMENT — WHAT A TINY FRACTION OF COMMITS WOULD DO

The French cloud market represents several billion euros per year. A fraction of these budgets — redirected to local providers — would allow them to fund R&D at the scale of their ambitions. To recruit the bare-metal engineers that large French companies let go. To build competitive managed services. To invest in datacentres powered by French renewable energy. The choice to spend with AWS or with OVH is a political choice as much as a technical one.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE DEALERSHIP — THE AUTOMOTIVE PARABLE

YOU WERE THE OWNER · YOU ARE NOW THE TENANT · YOU CALL THIS A STRATEGY

"Cloud strategy in three steps."

◆ STEP 1 — CONVINCE

Let go of your bare-metal engineers. "It is legacy." Close your cold aisle. "It is CAPEX, cloud is OPEX, it is more flexible." Abandon your physical network. "SD-WAN is the future." Validate your team's AWS certifications with training budgets. "It is upskilling, it is good." Brilliant. You have eliminated any possibility of going back.

◆ STEP 2 — DEPEND

Migrate your applications to proprietary managed services. DynamoDB because it is simple. BigQuery because it is powerful. Lambda because it is convenient. Pub/Sub because everyone does it. Six months later, these dialects are everywhere in your code. Your developers are AWS certified. Your CIO gives presentations on "our cloud journey". Your CTO quotes the Gartner Magic Quadrant. Everything is fine.

◆ STEP 3 — PAY

Receive the egress fees invoice when you try to retrieve your data. Discover your Reserved Instances are "noncancellable". Learn the applicable jurisdiction is Californian by reading the Terms — which nobody had read. Receive notification of a 30% price increase with notice of... none (GCP §2.6: "at any time"). Watch your €10 million annual cloud budget finance AWS's new datacentre in Spain.
أَتَسْتَبْدِلُونَ الَّذِي هُوَ أَدْنَىٰ بِالَّذِي هُوَ خَيْرٌ

◆ THE EXIT — IT EXISTS, IT IS NOT SIMPLE, IT IS POSSIBLE

The Sovereign Interface (corpus study) allows decoupling application code from the hyperscaler dialect. The Gentle Exit allows progressive migration without a catastrophic D-Day. Local providers — Ecritel, DRI, OVH, Scaleway — can receive what returns. Bare-metal competence can be rebuilt. Slowly. Painfully. But it can be. The corpus documents the path.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE COLD AISLE — PROOF THAT MASTERY OF IRON IS POSSIBLE

Cold aisle · Property of Amine RAITI · Mastery of iron exists — it is sometimes called a personal cold aisle

◆ WHAT THIS PHOTO SAYS

Mastery of iron is not an abstraction from the Opération Dindon corpus. It is a cold aisle. It is racks. It is hand-done cabling, dimensioned power supplies, configured network cards, updated firmware. It is someone who knows what is inside, why it is there, and what happens when it fails at 3am.

AWS has thousands of cold aisles like this one — and they generate 70% of the global cloud market from them. The question is not whether bare-metal is possible. The photo proves it is. The question is why organisations that could have this prefer to pay more not to have it — under California jurisdiction.

أَتَسْتَبْدِلُونَ الَّذِي هُوَ أَدْنَىٰ بِالَّذِي هُوَ خَيْرٌ

Are you trading what is better for what is inferior?
— Al-Baqara, 2:61

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART V

WOMEN IN INFRASTRUCTURE

4 pieces

HUMAN

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE SILENCE OF THE ROOMS

Why women are not in SRE

Structural diagnosis and levers for change

◆ FOREWORD — A REQUEST FOR READING

This document analyses a structural imbalance — the near-total absence of women from infrastructure, systems and network roles — using the same method as the previous studies in this corpus. It makes no judgment on individuals, no accusation against any profession, and no claim about the superiority or inferiority of anyone. It identifies mechanisms that produce this absence, and proposes levers to correct them.

This imbalance is a problem of efficiency as much as a problem of equity. Both arguments are presented here. Neither needs the other to stand.

Amine RAITI · Infrastructure Architect & SRE · Former instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

HUMAN

I

SECTION I · THE OBSERVATION

THE CLASSROOM AS A DIAGNOSTIC TOOL

This observation is not an impression. It is a repeated finding, across several years, across several vocational training centres, across different cohorts. Engineering classes: close to 50/50 gender balance. Developer training: the same. Systems & Networking: two women out of twenty, consistently. Sometimes zero.

The consistency of this figure across multiple years and multiple institutions rules out coincidence. It is not an atypical cohort. It is a pattern. And a pattern that repeats under the same conditions points to structural causes, not random ones.

◆ WHAT THE NUMBER SAYS AND DOES NOT SAY

Two women out of twenty in a Systems & Networking class does not mean women are incapable of doing this job. It means that something, upstream of competence, produces this absence at the stage of course selection and enrolment. The problem is not inside the classroom. It is in what happens before entering it.

The contrast with other technical disciplines is striking and is itself informative. If the issue were "women do not go into technical fields", the deficit should be visible across all technical training. It is not. It is concentrated on infrastructure, systems and networking — and on the physical terrain associated with it: the datacentre.

◆ NASSIHA — METHODOLOGICAL NOTE

The entire study rests on this first-hand empirical observation. It is not statistically representative in the academic sense. It is representative of an extended field experience, in a specific context. The mechanisms identified in the rest of this document are intended to explain this observation — not to prove it, as it is already established as a directly observed fact.

HUMAIN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE PHYSICAL EXPLANATION

GENUINE CAUSE OR CONVENIENT EXCUSE?

The first spontaneous explanation one hears — and is tempted to offer — is the physical one. The datacentre is a hot, noisy, physically demanding environment. A single intervention in a server room covers roughly 7,500 steps. Racks are heavy. Cables resist. There is something in this work that resembles a building site more than an office.

This explanation has an apparent logic. But it collapses against a simple, documented counter-example.

© Reuters — Women of the Moroccan Atlas

When construction machinery could not reach the villages of the Moroccan Atlas, it was women who carried the concrete blocks on their backs to build their village mosques. Physical resilience is not a genetically masculine trait. It is contextual, cultural, motivational. It is deployed where circumstances call for it.

◆ WHAT THE PHYSICAL EXPLANATION CONCEALS

A datacentre is no more physically demanding than a construction site, a professional kitchen, or a surgical suite. Each of those environments counts women among its practitioners. What makes the difference is therefore not the objective physical demand. It is the *representation* of that physical demand in the imagination of training institutions and recruiters — a representation that is constructed, and therefore changeable.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · STRUCTURAL MECHANISMS

THREE FORCES THAT PRODUCE THE ABSENCE

If the physical explanation does not hold, three structural mechanisms — each distinct, cumulative in their effects — can explain the observed imbalance. None of these mechanisms implies malicious intent on the part of individuals. They are system-level effects.

◆ MECHANISM 1 — THE CULTURAL REPRESENTATION OF INFRASTRUCTURE

The mental image of the SRE, the system administrator, the network engineer is masculine in popular culture — in television series, in conference photographs, in the illustrations of training materials. This image is constructed. It precedes young people's course choices by several years. A young woman choosing a career path is partly choosing a future image of herself — and that image tells her that server rooms are populated by men.

◆ MECHANISM 2 — THE ABSENCE OF FEMALE ROLE MODELS IN INFRASTRUCTURE

Role models reproduce through visibility. A sector that counts few women produces few visible female role models, which discourages new entrants, which keeps women scarce in the sector, which produces few role models — a loop closed on itself. This loop does not break spontaneously. It requires deliberate intervention at the level of training, communication and recruitment.

◆ MECHANISM 3 — THE IMPLICIT CULTURE OF TECHNICAL TRAINING

Infrastructure training materials are mostly written, illustrated and delivered by men, with examples, metaphors and a register that are often implicitly masculine. This is not conscious hostility — it is a blind spot. An instructor who does not think about it reproduces the context in which they were trained. Neutrality is not the result of a decision: it is the default effect of a non-decision.

◆ NASSIHA — WHAT THESE MECHANISMS DO NOT IMPLY

These three mechanisms are observable and correctable phenomena. They imply no claim about any intrinsic incapacity, no superiority or inferiority of any gender, no individual fault on the part of specific instructors or recruiters. Identifying a structural mechanism is not designating a culprit.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE TRAINING LOCK

TRAINING AS THE PLACE WHERE THE DECISION IS MADE

In "Anatomy of the Loss", training was identified as the central lever for reconquering infrastructure competence — against the dissolution of hard limits and the cognitive lock-in of the sector. The same logic applies here: if the gender imbalance is decided at the training stage, then training is where the levers for correction are found.

The shortage of women among SysAdmins and SREs is not decided at the point of hiring. It is decided at the point of orientation, and then consolidated during training. A young woman who does not walk into a Systems & Networking classroom will never have the opportunity to demonstrate that she can excel in this field. Discrimination through image precedes discrimination through competence.

◆ THE FOUNDATION OF IRON HAS NO GENDER

The pedagogical programme developed in this corpus — 26 weeks, from electricity to networking — was designed without any gender reference. Electricity is not masculine. Calculating subnets in binary is not masculine. Configuring an Active Directory domain controller is not masculine. These are technical skills, acquired through practice, rigour and repetition. They require no physical strength, and no genetic predisposition of any kind.

The iron will they demand — to understand, to persist through a fault that resists, to redo a cabling run three times until it is clean — is a human will. Not a masculine one.

The false perception of physical difficulty, described in section 2, acts as an invisible filter at the entrance to training. It is never written down. It is never stated explicitly. It is conveyed by the image of the room, by the instructor's language, by the composition of the cohort — and by the absence, within that cohort, of women who could have demonstrated through their presence that this perception is false.

◆ NASSIHA — TRAINING IS NOT THE ONLY LEVER

Training is the most accessible and most immediately actionable lever. It is not the only one. School guidance upstream, training centre communications, partnerships with secondary schools, and the visibility of women already working in the sector are complementary levers — acting over a longer timescale. This document focuses on training because that is where the author has direct experience and concrete proposals.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE COST TO THE SECTOR

A SECTOR RECRUITING FROM HALF ITS TALENT POOL

The systems and network infrastructure sector suffers from a shortage of qualified profiles that is documented, persistent and structural. Positions remain unfilled. Teams are chronically understaffed. The load per SRE engineer increases. On-call schedules grow denser. Burnout sets in.

In this context, effectively excluding half the working population from the talent pool is not a neutral choice. It is a choice that actively worsens an already existing shortage. A recruiter who does not think to recruit women, a training centre that does not think to attract women, a sector that does not think to retain them: each link in this chain halves the available talent pool. The market suffers directly.

> **◆ THE EFFICIENCY ARGUMENT IS INDEPENDENT OF THE EQUITY ARGUMENT**
>
> One does not need to be convinced by equity or equal opportunity arguments to recognise that recruiting from 50% of the available pool is a poor talent sourcing strategy. These two arguments — one moral, one economic — reach the same conclusion by different paths. The sector stands to benefit from hearing both of them.

There is also a quality argument. Diversity of profiles within a technical team produces diversity of approaches to problems. A datacentre seen through different perspectives is a datacentre whose blind spots are more often identified. Diversity is not a symbolic objective — it is a measurable operational advantage in terms of system resilience and diagnostic quality.

◆ WHAT INFRASTRUCTURE CONCRETELY LOSES

For every cohort of twenty Systems & Networking trainees that includes eighteen men and two women, the sector potentially loses ten additional qualified profiles who could have been trained, certified and operational. Over ten years, across hundreds of cohorts in dozens of training centres, this figure represents thousands of SRE engineers who do not exist because they were never trained. The shortage the sector deplores is partly of its own making.

◆ NASSIHA — THE LIMIT OF THE ECONOMIC ARGUMENT

The economic efficiency argument is the most immediately audible in a professional context. It also has a limit: it subordinates inclusion to economic usefulness, which is a fragile position if market conditions were to change. The equity argument — every person should have access to the roles for which they are qualified, regardless of gender — is more robust over the long term. Both are presented here, and both deserve to be heard.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · CONCRETE LEVERS

WHAT CAN BE DONE, STARTING NOW

The levers that follow are not grand statements of intent. They are precise actions, achievable within an ordinary training centre, without exceptional budget, by instructors who decide to think about this problem rather than not thinking about it.

◆ LEVER 1 — REPRESENTATION IN TRAINING MATERIALS

The photos used in course materials, the characters in lab scenarios, the examples cited in class — all of these imperceptibly construct the mental image of the profession. A material that systematically illustrates network administrators with masculine silhouettes or first names sends a signal that no one consciously formulated but that everyone receives. This signal is corrected by a simple decision: deliberately choosing mixed representations.

◆ LEVER 2 — VISIBILITY OF FEMALE INSTRUCTORS AND PROFESSIONALS

Inviting women SREs or system administrators to speak in training sessions — even for an hour, even via video call — breaks the visual monopoly of the masculine figure in the profession's imaginary. The role model who resembles who you are is the most powerful role model. It costs nothing more than a phone call to a professional willing to speak.

◆ LEVER 3 — LANGUAGE AND REGISTER IN THE CLASSROOM

The sporting, military or mechanical metaphors that sometimes pepper infrastructure courses are not neutral choices. They signal an implicit cultural context. They can be replaced by equally precise metaphors that are culturally broader, with no loss of pedagogical quality and a genuine gain in accessibility.

◆ LEVER 4 — MENTORING IN VOCATIONAL TRAINING CENTRES

Vocational training centres for jobseekers — the type from which the Foundation of Iron emerged — serve career-changers. A woman of 30 in a professional transition does not face the same barriers as a 17-year-old secondary student. She brings experience, maturity and motivation that can be considerable assets in a field where rigour matters more than age of entry. These centres are an underused talent pool.

◆◆◆

Infrastructure is the foundation of our digital ecosystem. There is no reason for it to be a private club. Not for ideological reasons. For reasons of competence, talent pool, and the long-term survival of a sector that needs every available talent.

◆◆◆

NEMO SUPRA LEGEM EST

HUMAIN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE INVISIBLE AMPUTATION

Why women were pushed out of bare-metal

◆ CONTEXT OF THE STUDY

This study extends "The Silence of the Rooms" — which documented women's absence from SRE — by posing a question that first study did not explicitly formulate: what if this absence was not accidental? It defends a three-level thesis — documented, structural, and hypothetical but coherent — on the relationship between women's exclusion from bare-metal infrastructure and the economic interests of hyperscalers. It concludes with the only structural response possible: early education.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

HUMAN

I

SECTION I · THE FIELD SRE — STRUCTURAL ENEMY OF HYPERSCALERS

THE ONE WHO CAN ALWAYS FALL BACK

The SRE engineer who masters infrastructure from the physical layer to the application layer is structurally hyperscalers' number one enemy. Not because they are hostile to them — but because they possess a capability that cannot be taken from them: the ability to fall back. To return to bare metal. To rebuild a sovereign infrastructure. To set the hard limits that cloud contracts never set.

This competence is an existential threat to the hyperscaler economic model. A client who can leave is a client who negotiates. A client who cannot leave is a captive. The lock-in strategy documented in "Anatomy of the Loss" aims to produce captive clients. And the most dangerous professional body for this strategy is precisely the SRE and systems administrator community who know how to build outside the cloud.

◆ WHAT A BODY AMPUTATED BY HALF LOSES SIMULTANEOUSLY

If this body is reduced by 50% through the systematic exclusion of half the population, it loses three distinct things simultaneously.

Numbers: the documented shortage of infrastructure profiles is directly amplified. Fewer engineers capable of falling back means fewer clients capable of negotiating.

A different perspective: women entering technical fields bring statistically a more conservative relationship to risk, a tendency to question implicit assumptions, and stronger resistance to technological fashion effects. This perspective is precisely the one that sets hard limits — and precisely the one that "soft on hard limits" DevOps seeks to replace.

Institutional legitimacy: a 95% male professional body is more easily caricatured, marginalised, and replaced by a generic title ("DevOps engineer") that can be positioned as more modern, more inclusive, more open — while being less competent on the physical layer.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · WHAT THE DOCUMENTATION SHOWS

THE NAMES IN THE CODE EXAMPLES

A thesis about deliberate exclusion deserves at least partial evidence. It will not be found in an internal memo — it is found in the configuration examples that thousands of engineers read every day as they learn their trade. AWS IAM documentation was examined.

◆ WHAT AWS IAM DOCUMENTATION SHOWS

Official IAM examples use a limited set of first names. Female names appear: Jane, Alice, Adele, Patricia, Chloe. Male names appear: John, Bob, David, Jim, Chris, Eli. Apparently balanced — until contexts are examined.

Female names appear systematically in examples of *basic user creation* and *restricted permissions*. Male names dominate examples of *administrator roles*, *AssumeRole*, *advanced permissions*, *role sessions in CloudTrail*.

The canonical example of an IAM security problem features "John Doe" creating the risk — "Jane Doe" inheriting it passively. The woman as a dependent variable of the man's behaviour, in the code example the beginner engineer reads on their first day of AWS training.

◆ NASSIHA — WHAT THIS SIGNAL PROVES AND DOES NOT PROVE

This documentation bias does not prove deliberate exclusion intent. It proves normalisation — the engineer who learns to "assume an admin role" learns that the person doing it is called John. A thousand engineers learning this form a mental representation. This mental representation informs hiring, promotions, and team cultures. This is not a conspiracy. It may be more effective than a conspiracy.

◆ HR LOOKING FOR "SOFT ON HARD LIMITS" DEVOPS

While documentation normalises the male profile in advanced administration roles, HR departments seek "DevOps engineers" — a title that, as documented in "The Infrational Crisis", dilutes real operational competences. This DevOps profile is presented as more inclusive, more collaborative, more modern than SRE or systems administrator. It is also structurally less dangerous for hyperscalers — because it does not master the physical layer.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE STRONG THESIS

WHAT IF IT WAS NOT AN ACCIDENT?

This section defends a thesis the author explicitly presents as a hypothesis — not a certainty. It is presented as such because it deserves to be thought through, even without direct proof. A thesis coherent with the interests at stake and with observable facts is not a conspiracy theory. It is a structural hypothesis.

◆ THE STRUCTURAL HYPOTHESIS

Hyperscalers have every interest in keeping the body of SREs and systems administrators — those who can fall back to bare metal — as small as possible, as lacking in legitimacy as possible, and as invisible as possible in public debates about the future of digital infrastructure.

The exclusion of women from this body — through the hostile culture documented in "The Silence of the Rooms", through masculine normalisation in documentation, through the replacement of SRE by DevOps in HR job descriptions — produces exactly this result. The body is amputated by half. It becomes less complete, less diverse in its perspectives, less capable of collectively resisting hyperscaler cognitive capture.

This exclusion did not need to be planned to have been produced. Economic systems produce outcomes that serve their dominant interests — often without anyone having explicitly decided to produce those outcomes. The question is not "was there a secret meeting?" The question is "who benefits from this outcome?" The answer is clear.

◆ THE CONVERGENCE OF MECHANISMS

Four mechanisms converge toward the same result without requiring explicit coordination:

— The hostile culture of bare-metal infrastructure teams (documented)
— Masculine normalisation in technical documentation (documented)
— The replacement of SRE by DevOps in HR job titles (documented)
— The devaluation of hard limits as a necessary competence (documented in "The Infrational Crisis")

Four distinct mechanisms, four distinct actors, one unique result: the professional body most dangerous to hyperscalers is amputated by half and marginalised in the public debate on digital infrastructure.

◆ NASSIHA — WHY DEFEND THIS THESIS DESPITE ITS UNPROVABILITY

A thesis does not need to be proven to be useful. It must be coherent with observable facts and point toward corrective actions that would be relevant even if the thesis were false. If women were accidentally excluded, correcting this is just. If they were deliberately excluded, correcting this is as effective as it is urgent. The strong thesis and the weak thesis lead to the same action.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE DIFFERENCE THAT DISTURBS

The diversity argument is often presented as a moral argument — we must include women because it is just. It is just. But it is also an operational argument that homogeneous infrastructure teams tend to ignore.

◆ THE RELATIONSHIP TO RISK

Organisational psychology and risk management studies converge on one finding: mixed groups make better-calibrated risk decisions than homogeneous male groups. Not because women are "less courageous" — but because they statistically have less tendency to follow group effects (groupthink) that push homogeneous teams toward overly aggressive decisions. In infrastructure, refusing groupthink has a name: hard limits. The engineer who says "no, we do not migrate this critical workload to this cloud service before we have an exit plan" is the engineer hyperscalers do not want in the room.

◆ RESISTANCE TO TECHNOLOGICAL FASHION EFFECTS

The culture of homogeneous male technical teams is particularly vulnerable to fashion effects — the shiny technology, the quarter's framework, the architecture that fills conferences. This vulnerability is documented and exploited. Hyperscalers have built conferences (re:Invent, Google Next, Microsoft Build) that are machines for producing fashion effects favourable to their services. A more diverse professional body is structurally more resistant to these fashion effects — because diversity of perspective breaks the group dynamics that fuel them.

◆ THE QUESTION NOBODY DARES ASK IN THE MEETING

In a homogeneous team under conformity pressure, the question "but can we do without this cloud service if the relationship with the provider deteriorates?" is hard to ask. It goes against the collective enthusiasm. Mixed teams ask these foundational questions more readily — because diversity of perspective legitimises the questioning of implicit assumptions. This questioning is exactly what the Opération Dindon corpus calls hard limits. And it is exactly what dominant DevOps culture discourages.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE DOCUMENTED EXCLUSION MECHANISMS

HOW YOU PUSH WITHOUT PUSHING

The exclusion of women from bare-metal does not pass through explicit discrimination decisions. It passes through an accumulation of mechanisms that make entry difficult, staying exhausting, and leaving natural. Each mechanism in isolation is defensible. The accumulation is systemic.

◆ MECHANISM 1 — ON-CALL CULTURE AS SELECTION

Nocturnal on-call in infrastructure is presented as a neutral professional requirement. It is not neutral. In a society where domestic and parental load still falls predominantly on women, uncompensated and unstructured nocturnal on-call is a gendered filter. Women do not "leave" bare-metal because they do not want to. They leave because the infrastructure operational model was not designed to be compatible with a life outside the office. "The Pager and the Body" documented this mechanism without naming its gendered dimension. This study names it.

◆ MECHANISM 2 — THE JOB TITLES THAT ERASE

The progressive replacement of "systems administrator" or "SRE engineer" by "DevOps engineer" in job postings is not trivial for women. The "DevOps" title appeared in a recruitment context that explicitly valued the collaborative, cross-functional profile, less anchored in the server room. This profile was presented as more inclusive — and in image terms, it is. But in terms of expected skills, it erases precisely the physical layer that women could have mastered as well as men. The profession was made "more welcoming" by removing its most strategically important content.

◆ MECHANISM 3 — THE ABSENCE OF ROLE MODELS IN DOCUMENTATION AND TRAINING

Section 2 documented the bias in AWS example names. The same bias exists in YouTube tutorials, Udemy courses, and the examples in technical reference books. When a young woman learns to configure a Linux server and all the examples show "john@server:~$" in the prompt, she receives a subliminal signal: this console is not for her. This signal is not a prohibition. It is friction. And accumulated frictions discourage without ever formally forbidding.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

ATTRACT GIRLS BEFORE THE FILTER PLAYS

All measures that seek to correct exclusion at the labour market entry point act too late. The decision not to pursue technical careers is made earlier — in middle school, in high school, sometimes in primary school. That is where the filter plays. That is where the correction must intervene.

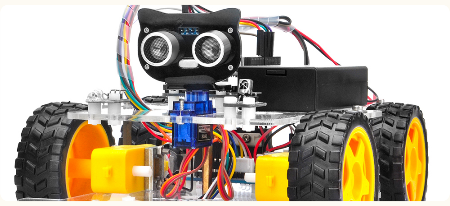

An Arduino project — the gateway into physical infrastructure

◆ ARDUINO, ROBOTICS, CONNECTED CARS — BEFORE THE BACCALAUREATE

Hands in the machine before the theory. An Arduino project that blinks an LED, then controls a robotic arm, then connects two microcontrollers — this three-year journey in middle school builds a hardware intuition that theoretical computer science classes do not build. The girl who assembled her own robot at 13 does not need anyone to explain what a GPIO, an I2C bus or an output voltage is. She lived it. And this lived experience changes the relationship to the physical layer for the rest of a professional life.

The connected car, the drone, home automation — these projects have an additional advantage: they are bridges between the physical world and the digital world. They show that infrastructure is not abstract. It is in objects. It is in the home. It is everywhere.

◆ A SINCERE CALL — WRITTEN BY MY OWN HAND

Our women — of every age — need to come and reinforce us in the face of abstraction. Bare-metal, cables, servers, the physical layer: this is ground where they have every right, every intelligence, every perspective. If they do not come, if we do not bring them, if we do not show them that this is for them as much as for us — it is going to be a mess. For infrastructure. For sovereignty. For all of us.

— AMINE RAITI

HUMAN

◆ WHAT COMPANIES CAN DO

Companies that want to correct the imbalance in their infrastructure teams do not have to wait for the market to do it. They can sponsor robotics clubs in nearby high schools. They can organise "infrastructure days" for final-year female students. They can mentor women in systems and networking HNC programmes. These actions cost little. They produce results over a 3-to-5-year cycle — not a quarter. They are the only diversity investments that work structurally.

◆ WHAT THE MEN IN PLACE CAN DO

The senior SRE who reads this study and recognises their own team in what it describes has a simple responsibility: make their team visible. Talk about their profession in schools. Invite female students to visit their server room. Share their own beginner struggles — the poorly connected cables, the failed configurations, the on-call nights. Show that this profession is learned, requires no particular gift, and is open to anyone who wants to practise it.

A professional body amputated by half cannot defend digital sovereignty at full strength. Reconstruction begins before the baccalaureate — not after hiring.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STUDY 65 · OPÉRATION DINDON · JUNE 2026 · CROSS-CUTTING STRATUM

IRON MIXITY

Breaking the Infrastructure Shortage — The Hidden Reservoir of Sovereign Digital

A Macro-Economic Audit of Recruitment Tension and an Activation Protocol for the Sovereign Long Apprenticeship

HUMAN

◆ THE THESIS — HUMAN CAPITAL AS A CROSS-CUTTING STRATUM

Twenty years of career — teacher, technician, system administrator, Head of SRE — converge on a single finding: the structural shortage in Iron professions is the mechanical consequence of amputating half the available human capital. None of the six strata of The Sovereign Exodus can produce its full effect without a seventh axis, cutting across the first six: the massive reconquest of female human capital in the low-level layers of infrastructure.

PAGES	RNCP TITLES	MECHANISM	WATERMARK	CORPUS REF.
16	2	PRO-A	HUMAN DOMINANT	65

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Trainer since 2006

CC BY-NC-SA 4.0 · *Opération Dindon · June 2026 · Claude/Gemini collaborative audit, 4 rounds*

THE IRON SHORTAGE IS NOT A DEMOGRAPHIC ACCIDENT — IT IS A MECHANICAL AMPUTATION

Twenty years of career — teacher, technician, system administrator, Head of SRE — converge on a single empirical finding: the structural shortage in Iron professions is not a demographic accident. It is the mechanical consequence of amputating half the available human capital.

When a technical field durably excludes 50% of its potential talent pool, it does not merely reduce its number of candidates. It cuts itself off from a diversity of approaches, profiles, retraining paths — and it locks itself into a spiral where the shortage itself becomes the argument justifying outsourcing to managed solutions, proprietary consoles, hyperscalers.

◆ THE CROSS-CUTTING STRATUM

None of the six strata of The Sovereign Exodus can produce its full effect without a seventh axis, cutting across the first six: the massive reconquest of female human capital in the low-level layers of infrastructure. This is not one stratum among others. It is what irrigates all the others.

This study documents three things: the real scale of the deficit, quantified and sourced; the structural parallel with other chronically tight sectors (metallurgy, construction); and the concrete protocol — regulatory, financial, pedagogical — to reverse the trajectory before the breaking point is reached.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · THE MACROSCOPIC DIAGNOSIS — VERIFIED FIGURES

3% OF WOMEN IN TELECOMS-NETWORKS — THE MOST AMPUTATED FIELD IN DIGITAL

These figures are not an activist estimate. They are INSEE and France Travail data, verified and sourced.

3%

of women in telecoms and network infrastructure jobs, versus 10% of men among digital professionals

INSEE, Employment surveys 2021-2022

24%

of women across all digital professions, versus 50% in other professions

INSEE, 2021-2023

In the training programmes that feed these fields, the gap widens before entry into the labour market: **17% OF GIRLS** in computer engineering programmes, **6%** in IT vocational tracks. The field does not lose its women along the way — it never recruits them.

WHAT THESE FIGURES CONCRETELY MEAN

A sector recruiting from just 3% of a potential pool of trained women does not have a "preference" problem — it has a capture problem. Tight-labour fields do not suffer from a lack of potential candidates. They suffer from a structural attractiveness deficit toward half the population.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

23% IN METALLURGY VERSUS 3% IN TELECOMS-NETWORKS — DIGITAL IRON IS THE MOST EXCLUDED FIELD

The parallel between digital infrastructure and physical transformation industries is not a circumstantial analogy. It is a structural homology, documented by the figures of the country's most exposed professional branch: metallurgy.

23%
of women in metallurgy's overall workforce
France Travail / UIMM, 2026

76,000
unfilled positions in the branch
France Travail / UIMM, 2026

This tension forced the French Union of Metallurgy Industries to launch, as early as 2023, the "*Tu as ta place*" ("You have your place") plan — six lines of action, an explicit and public objective: **MOVING FROM 23% TO 33% WOMEN IN THE INDUSTRY BY 2033**. This is not a CSR posture. It is an industrial war plan triggered by a shortage that threatens production.

◆ THE EXACT HOMOLOGY

Digital infrastructure suffers the same determinism — with an even sharper deficit. 3% of women in telecoms and networks versus 23% in metallurgy: digital Iron is the most amputated field of its potential pool, across all tight-labour branches combined. The digital Iron shortage will not be solved without activating the same talent-pool conquest lever that metallurgy had to trigger for its own industrial survival.

This parallel establishes one precise and verifiable thing: the durable exclusion of half the talent pool mechanically produces the same effect — a shortage that self-aggravates until a massive mixity plan corrects it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE CONCEPT OF "IRON MIXITY"

HYPERSCALERS DID NOT IGNORE MIXITY — THEY STEERED IT TOWARD THE UPPER LAYERS

◆ "WOMEN IN TECH" PROGRAMMES — A CAPTURE TOWARD GOVERNANCE

AWS, Azure and Google Cloud's "Women in Tech" programmes exist, are funded, and produce measurable surface-level results. But they systematically steer toward the upper layers of the cloud: governance, agility, Product Management, cloud consulting. Never toward the kernel, routing, physical cabling, low-level system administration.

THE CAPTURE MECHANISM

This is not a marketing coincidence. A hyperscaler that massively trains women in cloud governance and Product Management produces profiles that consolidate organisational dependency — not profiles that could audit or challenge it. Mixity, within this frame, never threatens the vendor's business model. It reinforces it.

◆ IRON MIXITY — TRAINING IN FUNDAMENTALS, NOT CONSOLES

Iron Mixity sets a different requirement: training women in raw engineering fundamentals — Bare-Metal, Linux, physical networks, system administration — rather than proprietary graphical consoles that perpetuate their employing organisation's dependency. Mastery of low-level layers is the only technical skill that is non-interchangeable and structurally immune to vendor lock-in. An engineer trained in the kernel and low-level network protocols depends on no proprietary ecosystem to practise her craft. A cloud console manager depends on it entirely.

◆ THE LINE OF DEFENCE

It is not cloud opponents who instrumentalise mixity. It is the hyperscalers who instrumentalise "Women in Tech" to turn them into brokers of their graphical consoles. Steering women exclusively toward cloud management is denying them access to the logical control of machines. Mixity without sovereignty is merely diversified subcontracting.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE ARCHITECTURE OF THE SOLUTION — THE TRIPTYCH

NEUTRAL FUNDAMENTALS · CRISIS MENTORSHIP · LONG APPRENTICESHIP

Three elements, and three only, structure the mechanism. None is sufficient alone.

1 NEUTRAL FUNDAMENTALS

Linux, physical networks, system administration — the technical base that makes a skill transferable from one environment to another, independent of any cloud vendor. This is the bedrock that distinguishes an Iron engineer from a console operator.

2 CRISIS MENTORSHIP

Senior system engineers, often sidelined by the NoOps waves documented elsewhere in this corpus, become transmitters of knowledge that disappears with them. Mentorship is not an HR option — it is the last transmission window before the know-how goes extinct.

3 LONG APPRENTICESHIP

A short training course is not enough to anchor raw engineering reflexes. The mechanism relies on a long apprenticeship — 12 to 24 months — combining theoretical training with genuine immersion in a production environment.

These three elements only work together. Training without mentorship reproduces the isolation documented in field dropouts. Mentorship without a contractual framework remains informal and unfundable. Apprenticeship without neutral fundamentals reproduces dependency on the tools of the vendor hosting the placement. The following sections detail the exact regulatory vehicle that makes this triptych fundable and enforceable.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE REGULATORY VEHICLE

THE SOVEREIGN LONG APPRENTICESHIP ALREADY EXISTS IN FRENCH LABOUR LAW

ARTICLE L6222-7-1 — FRENCH LABOUR CODE

The duration of an apprenticeship contract varies between six months and three years, set according to the type of profession and the level of qualification being prepared. This duration may be adapted by derogation, notably based on the candidate's initial skill level.

Two contractual vehicles coexist, with a determining age distinction:

APPRENTICESHIP CONTRACT

Capped at 29 years old, except for legal derogations (recognised disabled worker status, business creation/takeover project requiring the targeted skill).

PROFESSIONALISATION CONTRACT

No age limit for jobseekers aged 26 and over. This is the relevant vehicle for the pool of women retraining after 30.

◆ THE STRUCTURING CHOICE

For this study's target audience — women retraining professionally after 30 — the professionalisation contract is the appropriate legal lever. It requires no derogation, no exception, no heavy administrative process. It is built, by design, for exactly this profile.

The recommended 12-to-24-month duration fits strictly within the legal framework of both mechanisms. What remains to be established is the funding — the subject of the next section.

LEX

Amine RAITI · CC BY-NC-SA 4.0

RNCP36370 AND RNCP37680 — THE MOBILISABLE TITLES ARE ALREADY IN THE NATIONAL REGISTRY

The lock most often invoked against long-training mechanisms — the absence of an identified funding vehicle — does not exist here.

◆ RNCP36370 — SYSTEMS, NETWORK AND SECURITY ADMINISTRATOR

Registered title, accessible through an apprenticeship contract. Source: France Compétences.

◆ RNCP37680 — TP AIS (SECURE INFRASTRUCTURE ADMINISTRATOR)

Level 6 (Bachelor's equivalent), issued by the Ministry of Labour. Rhythm: 1 day at training centre / 4 days in company over 12 months — exactly the apprenticeship format proposed. 100% OPCO funding under apprenticeship.

The **OPCO ATLAS** fund publicly confirms funding for the "Systems and networks: Cloud systems and network administrator, Helpdesk technician, DevOps" track, with a public search engine by RNCP code giving the exact level of coverage.

INSERTION, NOT DIVERSION

The "Bare-Metal Fundamentals / Iron Mixity" module fits into the competence blocks already provided for by these titles — most systems/network curricula include an "infrastructure security and physical layer" block that can legitimately integrate this content. The choice of pedagogical support falls within the training centre's pedagogical engineering freedom, with no regulatory change required.

LEX

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE GUARDIAN TRANSMISSION PROTOCOL

MENTORSHIP AS AN ACT OF SURVIVAL FOR THE TECHNICAL CRAFT — NOT AN HR DIRECTIVE

◆ SENIORS — LAST CUSTODIANS OF VANISHING KNOWLEDGE

Senior system engineers — often sidelined or marginalised by the NoOps transformation waves documented elsewhere in this corpus — hold low-level knowledge that no longer transmits within current organisations. This knowledge disappears with every retirement, every layoff, every reorganisation that favours "cloud-native" over hardware mastery.

◆ THE NECESSARY PSYCHOLOGICAL REVERSAL

Presenting mentorship as a moralising HR directive produces rejection or surface-level compliance. Presenting mentorship as the only way to save know-how from programmed disappearance transforms the dynamic: the senior is no longer solicited in the name of diversity, but in the name of saving his own technical craft.

This reframing is not cosmetic. It changes the identity of the role: the senior becomes the last custodian of a rare skill, whose transmission is an act of professional preservation — not an additional burden imposed by a top-down diversity policy.

The transmission protocol structures this relationship: a formal pairing for the duration of the apprenticeship, precise and verifiable competence objectives, explicit recognition of the mentor role in the senior's own professional evaluation.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THREE PHASES OVER 24 MONTHS — PROTECTING THE APPRENTICE AND THE EXISTING TEAM

Integrating an apprentice into an SRE team under permanent strain carries a real risk: the overload that turns mentorship into a burden, and failure into premature abandonment.

M1-6 **SUPERVISED OBSERVATION**

Systematic pairing on interventions, with no isolated production responsibility. The goal is acquiring vocabulary, reflexes, gestures — not yet autonomy.

M7-18 **PROGRESSIVE RESPONSIBILITY**

Autonomous interventions on delimited, non-critical scopes, with after-the-fact validation by the mentor. Skill growth measured by precise indicators.

M19-24 **SUPERVISED AUTONOMY**

Taking on on-call duties and critical interventions, with the mentor on second-level standby. Transition to full-fledged engineer status.

This structuring protects both the apprentice — never exposed alone to a critical outage without a safety net — and the existing team, whose mentoring load is calibrated and decreasing over time rather than constant.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE HR REWRITING NUDGE

NEUTRALISING ALGORITHMIC ATS FILTERING THROUGH AGNOSTIC NOMENCLATURE

A perfectly designed training and mentorship mechanism fails if the downstream labour market continues to algorithmically filter the profiles it produces. Automated CV screening systems (ATS) are calibrated on proprietary terminology — "AWS Certified", "Azure Administrator" — which mechanically eliminates profiles trained on neutral fundamentals, regardless of their actual competence.

ALGORITHMIC NEUTRALISATION

A candidate certified RNCP36370 or holding the TP AIS, trained on Bare-Metal and pure Linux, can be eliminated before any human interview by an ATS filter configured on proprietary cloud keywords. The most rigorous training mechanism is useless if the job posting that precedes it is written in the vendor's newspeak.

The operational answer: systematically replace proprietary terminology with agnostic nomenclature in job postings — "Systems and Network Administrator" rather than "AWS Cloud Engineer", "Infrastructure Engineer" rather than "Azure DevOps Engineer". This reframing relies directly on the RNCP codes themselves, which are, by design, vendor-agnostic.

This nudge requires no negotiation with ATS vendors. It acts upstream, on the very wording of job offers — a lever entirely controlled by the hiring company.

LEX

Amine RAITI · CC BY-NC-SA 4.0

FROM ISOLATED GESTURE TO ENFORCEABLE STANDARD — SOCIAL DIALOGUE AS LEVER

The HR rewriting nudge, deployed at the scale of a single company, remains an isolated gesture. Deployed at the scale of a professional branch, it becomes an enforceable standard. Collective bargaining agreements and job classifications defined by professional branches carry a normative force that the internal practices of a single isolated company lack.

◆ THE ENFORCEABILITY MODEL

A professional branch that registers "Systems and Network Administrator (all infrastructures)" as the reference title, rather than implicitly validating segmentation by cloud vendor, creates a precedent that ATS systems and recruiters cannot ignore without exposing themselves to non-compliance with their own collective agreement.

This lever goes beyond the scope of an individual study. It calls for social dialogue work — with digital employer organisations (Numeum) and representative trade unions — to have agnostic nomenclature recognised as a branch standard, not as an isolated and fragile initiative.

LEX

Amine RAITI · CC BY-NC-SA 4.0

TURNING THE COMMERCIAL APPARATUS OF EX-SSII TOWARD NEUTRAL SKILLS

IT consultancies (ESN) — formerly known as SSII — employ a significant share of France's system and network engineers. Their commercial apparatus, today oriented toward reselling hyperscaler-certified skills, can be turned around. An ESN that builds a pool of profiles certified in neutral fundamentals — LFCS (Linux Foundation Certified Sysadmin), RHCSA (Red Hat Certified System Administrator) — holds a differentiating commercial argument against a public procurement market increasingly sensitive to digital sovereignty issues.

THE PUBLIC PROCUREMENT LEVER

A public tender that explicitly requires agnostic certified skills — rather than proprietary certifications from a named cloud vendor — mechanically steers bidding ESNs toward building this talent pool. Public procurement becomes an instrument for structuring the technical labour market, with no new regulatory constraint.

This mechanism connects directly with the training scheme described in the previous sections: apprentices trained on RNCP36370 or TP AIS are precisely the pool that ESNs, under public procurement pressure, will have an interest in recruiting and showcasing.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

CAPTURING THE TALENT POOL CANNOT WAIT TEN YEARS

Three factors converge to create a tightened time window:

◆ THE MASS DEPARTURE OF SENIORS

Every year that passes without organised transmission is a year of know-how lost irreversibly, with no recourse once the departure has occurred.

◆ THE ACCELERATION OF HYPERSCALER CAPTURE

"Women in Tech" programmes steered toward the upper layers of the cloud continue to capture a growing share of the available female talent pool, shrinking the basin mobilisable for Iron Mixity by the same amount.

◆ THE BREAKING POINT OF SOVEREIGN INFRASTRUCTURES

The mechanisms documented in The Sovereign Exodus (Bare-Metal, sovereign IDP, cloud-free BCP) require trained workforces to be operational. Without renewing the talent pool, these infrastructures will remain theoretical for lack of hands to operate them.

◆ THE URGENCY IS NOT RHETORICAL

An activation plan launched now produces its first operational graduates within 12 to 24 months — exactly the timeline on which the credibility of the entire Opération Dindon corpus is staked. Deferring this launch by five years means facing the breaking point without the human capital needed to confront it.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

EVERY CELL REFERS TO AN EXISTING MECHANISM — NOTHING REQUIRES PRIOR REFORM

DIMENSION	MECHANISM	REFERENCE
Target audience	Women retraining professionally, aged 26 and over	*No upper age limit*
Contractual vehicle	Professionalisation Contract	*No age limit for 26+ jobseekers*
Legal basis	Duration 6 months to 3 years, adaptable	*Article L6222-7-1, French Labour Code*
RNCP title #1	Systems, network and security administrator	RNCP36370
RNCP title #2	Secure Infrastructure Administrator (TP AIS)	RNCP37680, *Level 6, Ministry of Labour*
Funding	100% coverage under apprenticeship	OPCO Atlas
Recommended duration	12 to 24 months, in 3 phases	*Observation · Responsibility · Autonomy*
Shortage indicator	3% women (telecoms-networks) vs 23% (metallurgy)	INSEE 2021-2022 · *France Travail/UIMM 2026*
Reference target	33% mixity by 2033 (metallurgy model)	"*Tu as ta place*" *plan*, UIMM, 2023

◆ WHAT THIS TABLE DEMONSTRATES

Every cell of this table refers to an existing mechanism, a verified source, an active regulatory code. Nothing in this architecture requires prior legislative reform. The protocol is executable now, with the legal and financial tools already in place.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WITHOUT WOMEN RETURNING TO IRON, THE SOVEREIGN EXODUS REMAINS A PLAN WITHOUT HANDS

The INSEE and France Travail figures cited in this study are verified and sourced. One additional data point deserves to be named for what it is: an empirical estimate, not an official statistic.

AMINE RAITI'S ESTIMATE

Over twenty years of career — engineering school teacher, technician, system administrator, Head of SRE — the field estimate is that roughly 90% of the potential female talent pool for low-level infrastructure roles remains untapped. This figure is not an official statistic. It is a direct observation, repeated over two decades, in real production environments.

This distinction between verified data and field observation does not weaken this study's thesis. It strengthens it: official figures already suffice to demonstrate a major structural deficit. The empirical estimate adds a dimension that aggregated statistics do not always capture — the conviction of a practitioner who has watched, over two decades, the talent pool empty out without ever refilling.

IRON MIXITY IS NOT ONE STRATUM AMONG SIX

It is what allows the other six to produce their full effect.
Without the massive return of women to the low-level layers of infrastructure,
the Sovereign Exodus remains a plan with no hands to execute it.

AMINE RAITI — INFRASTRUCTURE ARCHITECT & SRE

Former engineering school professor · Trainer since 2006

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

MATERNITY AS CONTINUITY

A PROPOSAL TO ENSURE THAT BIRTH IS NOT A CAREER INTERRUPTION

HUMAN

◆ **CONTEXT OF THE STUDY**

This study is part of the Opération Dindon corpus, following "The Silence of the Rooms" (women's absence from SRE) and "The CPF Misused" (training funding). It addresses an angle those studies did not cover: not the entry of women into infrastructure roles, but their retention at the moment of maternity. It puts forward a concrete proposal built on three principles — maintaining the link, full financial protection, and anti-downgrade by design — without creating any new legal mechanism.

◆ **PRELIMINARY NOTE**

This proposal recognises maternity as a biological function from which society as a whole benefits, and whose career cost is today borne disproportionately by the women who carry it out. It prescribes no behaviour — it creates conditions in which every woman can freely choose the form of her maternity without suffering professional consequences for it.

AMINE RAITI · Infrastructure Architect & SRE · Former instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE TWO PROBLEMS

WHAT MATERNITY COSTS TODAY

Maternity is a biological function. It is not a lifestyle choice, a career decision, or a voluntary absence. It is a reproductive act from which society as a whole benefits — the future workers, the future taxpayers, the future engineers are children that someone carried, delivered and fed. This elementary fact is systematically ignored in the current organisation of work, which treats maternity as an interruption to be managed rather than a contribution to be recognised.
In systems and network infrastructure roles, this ignorance produces two concrete, documented problems.

◆ PROBLEM 1 — KNOWLEDGE LOSS IN THE TEAM

An SRE who goes on maternity leave takes with her a knowledge of the infrastructure that no one else in the team possesses in exactly the same way. She knows the failure paths, the architecture decisions made in a meeting three years ago, the reasons why a given configuration exists rather than another. For 4 to 16 months, the team operates without this knowledge. When a fault occurs at 3am, the intervention that would have taken 12 minutes takes 4 hours — or does not happen at all. This is not a sentiment problem. It is a measurable operational problem.

◆ PROBLEM 2 — THE CAREER PAUSE AND THE RETURN RISK

While feeding the future SRE, the SRE Girl watches her career pause. Certifications evolve. The infrastructure changes. Teams reorganise. On return, the risks are multiple: role downgrade under cover of "reorganisation", loss of knowledge about an infrastructure that evolved without her, misalignment with a team that has formed new habits. These risks are documented, frequent, and rarely sanctioned because they are hard to prove.

◆ NASSIHA — WHAT THIS STUDY DOES NOT CLAIM

This study does not claim that all companies mistreat employees on maternity leave. It identifies structural mechanisms that produce these risks independently of individual goodwill. A well-meaning manager in a poorly designed organisation produces the same outcomes as an indifferent one. The solution must be structural, not personal.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE FOUNDING PRINCIPLE

ANTI-DOWNGRADE BY DESIGN

Most existing proposals to protect women on return from maternity leave try to fix the problem after it has occurred — legal prohibition of downgrading, employment tribunal recourse, mandatory retraining. These approaches are reactive. They assume a rupture has taken place, and seek to repair it.

The proposal in this study is different in nature: it seeks to eliminate the rupture itself. A woman who never left her team cannot be downgraded on her return. There was no return. There was a continuation at reduced intensity.

◆ THE PRINCIPLE: MAINTAINING THE LINK, NOT REAL PRODUCTION

The distinction is fundamental. Maintaining the link is not a work demand. It is the possibility, at the exclusive initiative of the woman, to listen to a Teams call, follow a reorganisation, be aware of an architecture decision. 1 to 2 hours a day, when she decides, if she decides. No deliverable, no objective, no evaluation. The baby always comes first — this hour only exists if and only if the woman chooses to take it that day.

The counterpart of this maintained link is twofold: the team preserves its human expertise, and the woman stays in the flow of the team. She follows the reorganisations. She knows the new decisions. She sees who has taken which role. When she gradually increases her time, she has nothing to catch up on — she was there.

◆ THE HARD LIMIT AGAINST THE DOUBLE EDGE

This mechanism has one identified risk: the employer who turns the option into an obligation. The hard limit is simple and non-negotiable — **ALL PROFESSIONAL CONTACT DURING THE BIOLOGICAL RECOVERY PERIOD MUST BE INITIATED BY THE EMPLOYEE, NEVER BY THE EMPLOYER.** No requests, no implicit expectations, no evaluation based on presence or absence during these 1-2 hours. The woman is the sole decision-maker on when, whether, and how much.

◆ NASSIHA — THE BIOLOGICAL RECOVERY PERIOD IS UNCONDITIONAL

The proposal distinguishes two strictly separate phases. The biological recovery phase — duration set by the woman, professional contact forbidden — is unconditional and non-negotiable. It functions like a standard leave: the employer has no access. The link-maintenance phase only begins when the woman decides she is ready, and lasts only as long as she chooses each day.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · FULL FINANCIAL PROTECTION

ZERO LOSS, ZERO DOWNGRADE

Maintaining the professional link is not enough if the financial protection is incomplete. An SRE who stays connected to her team but loses 16% of her net salary, sees her supplementary pension frozen, and fails to accumulate the same rights as her male colleagues during the same period — this SRE suffers a double penalty: she works to maintain her expertise during her maternity, and she is still financially penalised for having lived it.

◆ COMPONENT 1 — EMPLOYER TOP-UP TO FULL NET SALARY

Social Security maternity benefit represents approximately 84% of capped net salary. The proposal requires the employer to top up to the full net salary — with no ceiling. This mechanism already exists in the Syntec collective agreement for sick leave beyond one year of seniority: the employer tops up the Social Security benefit to maintain full net salary. The same logic, applied to maternity, creates no new mechanism — it extends a proven one.

◆ COMPONENT 2 — PENSION CONTRIBUTIONS PAID BY THE STATE

Basic state pension contributions are partially covered during maternity leave via the CNAV. But supplementary pension contributions (AGIRC-ARRCO) and contributions in special schemes are not, or insufficiently. The proposal extends national coverage to all pension contributions — basic and supplementary — for the full duration of the leave. The justification is economic as much as moral: feeding a child in their first months is educational work from which society benefits. This work deserves compensation in rights, not a hole in the pension.

◆ COMPONENT 3 — SENIORITY AND SALARY PROGRESSION MAINTAINED

The maternity period is counted in full as seniority. No salary progression freeze. No missed appraisal that creates a gap with colleagues. The woman returns exactly to the level she would have reached had she not had a child. This is not an advantage — it is the neutralisation of a structural disadvantage.

◆ NASSIHA — THE SYNTEC MODEL AS LEGAL PRECEDENT

The Syntec agreement provides for full net salary maintenance during sick leave beyond one year of seniority. This precedent demonstrates that employer top-up to full salary is legally and economically viable in the digital and intellectual services sectors — precisely where SREs work. Extending this mechanism to maternity does not require inventing a new right. It requires applying an existing one to a situation the Syntec has not explicitly addressed.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · WHO IS COVERED

REMOTE-ELIGIBLE AND NON-REMOTE ROLES

The proposal addresses two distinct populations based on the nature of the role. Their treatment differs in mechanics, but is identical in rights.

◆ POPULATION 1 — REMOTE-ELIGIBLE ROLES

Every role performed remotely during the 2020-2021 COVID crisis is presumed remote-eligible. This precedent is objective, documented, and uncontestable: if the employer proved during 18 months that the role functioned at a distance, they can no longer claim that the same role cannot be performed from home during maternity. The burden of proof to the contrary belongs to the employer — not to the employee. For this population, the full mechanism applies: biological recovery phase, then optional link-maintenance at 1-2h/day at the woman's initiative.

◆ POPULATION 2 — NON-REMOTE ROLES

For women whose role objectively cannot be performed remotely — nurse, cashier, production operative — the link-maintenance mechanism via 1-2 daily hours is not applicable in its professional form. They receive the same financial rights by national solidarity: full net salary maintained, all pension contributions paid by the State, seniority and salary progression maintained. Without any presence counterpart — this is a right, not an exchange.

◆ FUNDING FOR NON-REMOTE ROLES: A BRANCH FUND

The additional cost of full protection for non-remote roles — the gap between Social Security benefit and full net salary — is funded by a pooled fund fed by the collective agreements of each industry sector. This pooling mechanism is structurally close to the provident funds already existing in most collective agreements. The number of births per sector is predictable, the cost is calculable, and the volume is not systemic: it is a one-time event per woman, not a permanent charge. Existing provident funds absorb far more volatile risks without difficulty.

◆ NASSIHA — SOLIDARITY IS NOT A FAVOUR

Funding the maternity rights of non-remote roles is not an act of generosity — it is the recognition that reproduction is collective social work, not a private individual decision whose costs must be borne by the person who carries it out. The same logic that justifies family allowances justifies full career rights protection during maternity.

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · WHAT IT RESOLVES IN THE TEAM

THE KNOWLEDGE THAT DOES NOT LEAVE

The proposal resolves both problems identified in section 1 simultaneously — and the resolution of the second is the direct consequence of maintaining the link, not a separate objective to be achieved by other means.

◆ THE TECHNICAL KNOWLEDGE THAT DOES NOT LEAVE THE TEAM

An SRE who maintains 1 to 2 hours of daily contact with her team stays informed of incidents, architecture decisions, and infrastructure evolutions. She is not available for on-call duties, emergency interventions, or long meetings. But she is reachable for a precise question that no one else in the team can answer. This partial, self-initiated availability is enough to keep alive the tacit knowledge she alone holds — those undocumented understandings that make the difference between a fault resolved in 12 minutes and a night-long incident.

◆ NATURAL KNOWLEDGE TRANSFER WITHOUT FORCED DOCUMENTATION

The current alternative to absence is the "handover" — documenting in a few weeks what the SRE has known for years. This documentation is always incomplete because tacit knowledge cannot be fully documented. Maintaining the link requires no documentation — it keeps the source alive. The knowledge remains available in the person who holds it, without having had to freeze it artificially in a document that no one will find at the right moment.

◆ ANTI-DOWNGRADE BY DESIGN — THE DEMONSTRATION

A woman who followed her team's reorganisations during her maternity does not return to an unfamiliar environment. She knows the new faces, the new roles, the new technical directions. She knows that service X was migrated, that team Y merged with Z, that the project starting when she left is now in production. She has nothing to catch up on — she was there, at reduced intensity, but there.

An employer seeking to downgrade this profile on her return would face a simple reality: this woman is more up to date than many of her colleagues who were on holiday, in training, or on assignment abroad during the same weeks. The downgrade has no technical argument left to invoke.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE COMPLETE PROPOSAL

THREE PHASES, ZERO RUPTURE

◆ PHASE 1 — BIOLOGICAL RECOVERY (DURATION SET BY THE WOMAN)

Content: full rest, professional contact forbidden except at the employee's own initiative.
Status: identical to standard leave — the employer has no access, no expectations.
Funding: Social Security benefit + employer top-up to full net salary (extended Syntec model) + full basic and supplementary pension contributions paid by the State + seniority and salary progression maintained.
Hard limit: any attempt at professional contact by the employer during this phase is an enforceable infringement — reversed burden of proof on the employer.

◆ PHASE 2 — MAINTAINING THE LINK (AT THE WOMAN'S INITIATIVE)

Content: 1 to 2 hours per day, decided by the woman, decided each day. Listening to a Teams call, following a reorganisation, a punctual question. No deliverable, no objective, no evaluation.
Status: optional, initiated exclusively by the employee. The employer can neither request, expect, nor take into account presence or absence during these hours in any subsequent appraisal.
Funding: hours performed are compensated at the pro-rata hourly rate, in addition to maternity rights — not in substitution.

◆ PHASE 3 — NATURAL PROGRESSIVE RETURN

Content: progressive increase of working time at the woman's initiative, up to full return.
Status: no imposed date, no minimum threshold. The woman builds her own ramp-up.
Guarantee: the role, classification, salary and team are identical to those at departure. The woman has nothing to catch up on because she never left.

Feeding the future SRE is work. Society benefits from it. Society must compensate it — in rights, in protection, and in career continuity for the person who carries it out.

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART VI

THE PEOPLE

8 pieces

HUMAN

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE YEMENI DINNER

What an evening at €78 per person does that a €2,000 offsite cannot

HUMAN

◆ CONTEXT OF THE STUDY

This study steps outside the technical perimeter of the Opération Dindon corpus to address a direct management angle: team cohesion in infrastructure. It starts from an observation as Head of SRE — that real team cohesion is not built in corporate offsite events, but in carefully chosen human moments. It illustrates this thesis with a concrete, costed, reproducible evening programme.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE PROGRAMME

AN EVENING FOR 5 — CHAMPS-ÉLYSÉES, CHÂTELET

The programme that follows is not a theoretical proposal. It is a programme that was executed, with an SRE team of 5 people. It is reproduced here as-is — with the timings, the locations, the total cost — because the precision of the concrete is more useful than the generality of a principle.

◆ THE EVENING PROGRAMME

18:00 — Leave the office together. No individual cars or metro journeys. The group leaves as a group.
18:45 — Bowling session. Walked from the office to the bowling alley — approximately 45 minutes on foot through Paris. The walk is not a commute: it is the first activity of the evening.
20:30 — Gastronomic dinner at a Yemeni restaurant on the Champs-Élysées. A deliberate choice — not a standard brasserie, not the usual company restaurant. A place that surprises, that has something to say, that opens a conversation beyond work.
21:00 — Move to Châtelet for a drink. A final informal moment, standing, before parting.
22:00 — Everyone goes home. Happy. Clear-headed.
TOTAL COST: €390 ALL-INCLUSIVE for 5 people, author's hotel night included.
COST PER PERSON: €78.

The shared dish — the team eating together by hand from the same plate, as in Morocco · © Amine RAITI

◆ NASSIHA — WHAT IS NOT IN THE BUDGET

€390 is the direct cost. What is not in that figure: no team-building consultant, no seminar room, no conference badge, no out-of-Paris travel costs, no difficult morning after. The true cost of a two-day corporate offsite for 5 people — travel, hotel, facilitator, room, meals — regularly exceeds €10,000. The ROI of the evening, measured in cohesion quality produced per euro spent, is not comparable.

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · WHY IT WORKS

THREE MECHANISMS THAT THE OFFSITE CANNOT REPLICATE

The evening is not effective despite its modest cost. It is effective partly because of it — because the absence of corporate overproduction leaves room for the real human mechanisms of connection. Three of them are at work in this programme.

◆ MECHANISM 1 — WALKING REMOVES HIERARCHY

45 minutes on foot through Paris, as a group, before any activity. This is not a commute — it is the first conversation of the evening. People talk differently when walking side by side than when sitting face to face in a meeting room. Walking removes the desk, the badge, the org chart. It puts everyone at the same pace, literally. The Head of SRE and the junior walk in step. The hierarchy does not disappear — it goes into standby for the evening.

◆ MECHANISM 2 — THE GAME REVEALS AND LIBERATES

Bowling creates a light competition, without professional stakes, in which anyone can shine or fail without consequence. The most reserved engineer in meetings might be the best player. The manager might throw a gutter ball to the laughter of the whole team. These momentary role reversals are valuable: they reveal personality facets that the professional context masks, and they create shared memories — "remember when you got a strike on the last throw" — that weave the fabric of the group over time.

◆ MECHANISM 3 — THE CHOICE OF VENUE SAYS SOMETHING ABOUT THE MANAGER

A Yemeni restaurant on the Champs-Élysées is not a neutral choice. It is not the brasserie next door, not the usual company restaurant. It is a choice that says: I am taking you somewhere you may not know, a cuisine you may never have tasted, a place that tells a story. This choice positions the manager as someone curious, open, off the beaten track — and this image reflects onto the team, which recognises itself in its manager.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE 22:00 CLOSE

THE DETAIL THAT CHANGES EVERYTHING

There is one element of the programme that most team-building events ignore, and it makes all the difference: the evening ends at 22:00. Everyone goes home happy, clear-headed. This is not a logistical constraint — it is a deliberate management decision.

◆ WHY 22:00 AND NOT MIDNIGHT

An evening that extends beyond a certain threshold produces risks the manager must anticipate: the fatigue that turns good humour into irritability, the alcohol that lifts inhibitions in the wrong direction, the subgroups that form and exclude, the words that exceed the thought and are remembered the following morning. "Ending on a good note" is not a cliché — it is a technique. People remember how an evening ends, not just how it begins.

◆ "HAPPY AND CLEAR-HEADED" — WHAT THIS MEANS FOR THE TEAM

The following morning, everyone arrives at the office with the same shared memory — a good evening, no incident, no regrets. This shared memory is a trust capital the team can spend in the weeks that follow: to approach a difficult topic with less rigidity, to trust a colleague on a point they might previously have challenged, to get through a 3am production incident with the certainty that they will come out the other side together.

◆ COHESION IS NOT A STATE — IT IS A REGULAR INVESTMENT

A single evening does not create team cohesion. It maintains it, reinforces it, releases the accumulated friction. Cohesion is like infrastructure: it degrades if not maintained. One dinner per quarter — €78 per person, 4 times a year — represents €312 of annual investment per person in team cohesion. A production incident poorly managed because two team members only half-trust each other costs far more than that in extended MTTR, suboptimal decisions, and unnecessary escalations.

◆ NASSIHA — THE MANAGER WHO STAYS IN A HOTEL

The programme notes a hotel night for the author — included in the €390. This detail says something important: the manager lives far enough away to need a hotel, but organises the evening in Paris for their team anyway. This is not heroic sacrifice — it is a clearly stated priority. The team sees it. They remember it.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · WHAT THE OFFSITE DOES NOT DO

THE PROBLEM WITH FORCED COHESION

The corporate team-building offsite is the instinctive reflex of organisations that want to improve cohesion. Two days in a hotel with a consultant, trust exercises, structured group activities, a presentation on personality styles. This format exists because it is billable, schedulable, and produces a report. Not because it is effective.

◆ FORCED COHESION PRODUCES ITS OPPOSITE

Putting 10 people in a seminar room for two days with the declared objective of "strengthening team cohesion" creates a social pressure that inhibits precisely what it seeks to produce. People know they are supposed to bond. This awareness produces a performance of bonding rather than genuine bonding. They play along with the exercises. They participate in the workshops. They go home two days later with offsite memories, not team memories.

◆ THE TIMESCALE PROBLEM

An annual two-day offsite represents 16 hours of "cohesion time" across 365 working days. That is 0.4% of working time dedicated to relationships between team members. The remaining 99.6% unfolds in meetings, incidents, tickets, and deadlines. Cohesion built over 16 intensive hours once a year is no more durable than cohesion maintained over 4 hours per quarter in a carefully chosen informal setting.

◆ WHAT ACTUALLY CREATES CONNECTION

Social psychology research on group cohesion converges on a few necessary conditions: sharing a common emotional experience, mutual discovery outside usual roles, and the feeling of having been chosen rather than summoned. The evening — walking, bowling, Yemeni dinner, drink at Châtelet — ticks all three. The offsite with its pre-formatted agenda and imposed participant list generally ticks none of them.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · COHESION IN SRE TEAMS

WHY IT IS MORE CRITICAL IN INFRASTRUCTURE

Team cohesion matters in all professional contexts. It is critical in infrastructure. The reason is simple: SRE teams make decisions under maximum stress — at 3am, with a production system going down, with users or clients waiting, with limited time to diagnose and act. In these conditions, the quality of the relationship between team members is not a comfort — it is an operational variable.

◆ TRUST UNDER INCIDENT

During a major incident, two trust dynamics are critical. Technical trust — "I know my colleague knows what they are doing" — and relational trust — "I know my colleague will be honest about what they do not know". The second is harder to build than the first, and it is built precisely in moments outside work. An engineer who has seen their manager throw a gutter ball and laugh about it knows that manager can acknowledge their limits. This link between the informal and the professional is real and documented.

◆ THE INTERNAL CONFLICTS THAT DISAPPEAR

Tensions in an infrastructure team accumulate in the daily routine — poorly handed-off tickets, contested architecture decisions, unevenly distributed on-call, meetings that overrun. These frictions are normal. They become toxic when they have no outlet. The evening is the outlet. Not because it erases the problems — but because it puts the people first and the problems second. You do not resolve conflicts over a Yemeni dinner. You remember that the people you work with are people, not nodes in an org chart.

◆ NASSIHA — COHESION DOES NOT REPLACE MANAGEMENT

A team that goes out together once a quarter but whose structural problems are not addressed — excessive on-call load, absent recognition, unclear direction — will not be saved by bowling. Cohesion is a multiplier: it amplifies what is already in place, for better or worse. A good team with cohesion is excellent. A dysfunctional team with cohesion is more loyal in its dysfunction. The evening does not replace management — it complements it.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

THE QUARTERLY DINNER PROTOCOL

The proposal of this study is concrete and reproducible. It requires no exceptional budget, no complex HR approval, no consultant. It requires a manager who decides that their team's cohesion is an operational priority and allocates 4 evenings per year to maintain it.

◆ **THE FORMAT PRINCIPLES**

Frequency: once per quarter. Often enough to maintain the bond, rare enough that the evening remains an anticipated event.
Size: ideally 4 to 7 people. Below that, the group effect does not occur. Above it, subgroups form naturally and the team does not experience the evening as a unit.
Three-part structure: a physical or playful activity (bowling, karting, escape room, walk) — a dinner in a thoughtfully chosen, non-formatted venue — a short informal moment before parting (a standing drink, not a second restaurant table).
Fixed closing time: 22:00 or 23:00 maximum. Define the end time before you start. End on a good note.
Target budget: €60 to €100 per person depending on the city. Less than that, you cut corners on the experience. More than that, you create a reciprocity obligation that weighs on those with financial constraints.
The venue choice is a statement: choose a restaurant that says something — a discovered cuisine, an unexpected neighbourhood, a place the manager loves and wants to share. Not the first TripAdvisor result.

◆ **WHAT €312 PER PERSON PER YEAR BUYS**

Four evenings at €78, that is €312 per person per year. For this budget, a manager buys: four 4-hour evenings where their team remembers it is a team. Production incidents handled with more fluidity because the trust exists before the crisis. Conflicts that dissolve because people see each other as people rather than roles. And reduced turnover — because you do not easily leave a team you have shared memories with.

◆◆◆

For not much money, you can maintain seamless team cohesion — the kind where internal conflicts disappear over a good dinner. The Yemeni restaurant on the Champs-Élysées proves it at €78 per person.

◆◆◆

NEMO SUPRA LEGEM EST

HUMAIN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE BODY FIRST

From the Barbès hammam to the Belleville dinner

The neurochemical sequence at €35 per person

HUMAN

◆ CONTEXT OF THE STUDY

This study extends "The Yemeni Dinner" and "Play in SRE Teams" on a new angle: physical proximity as cohesion infrastructure. It argues that a hammam + free walk + neighbourhood dinner evening produces a complete, documented neurochemical sequence — for €35 per person — that no corporate seminar reproduces at this price or at this depth of human connection.

Amine RAITI — Infrastructure Architect & SRE

Teaching since 2006 · Electricity to Kubernetes · All audiences

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE PROGRAMME — CONCRETE, COSTED, REPRODUCIBLE

BARBÈS 7PM — BELLEVILLE 9PM — EVERYONE HOME BY 11PM

This programme is not a theoretical proposal. It is a real itinerary, with locations, timings and a budget. Its strength lies in its precision and accessibility.

◆ THE EVENING PROGRAMME — PARIS, REPRODUCIBLE EXAMPLE

19:00 — Meeting point at the Barbès hammam. The neighbourhood hammam is rented entirely for the team — 1 hour, €150 to €200 for 10 people. No corporate badge, no logo on the towel. Heat, steam, silence or conversation as preferred. The body exits the production vigilance state.

20:00 — End of hammam. Everyone makes their own way to Belleville — walking alone, walking in a spontaneous small group, or by metro. No imposed route. Sub-groups form naturally based on the affinities of the moment, not the org chart. The geographic transition Barbès → Belleville is the progressive decompression between the rested body and the dinner.

21:00 — Meeting point communicated in advance: Asian restaurant in Belleville. Not a reservation in a corporate restaurant. A place chosen for its character — the food, the neighbourhood, the atmosphere. The team gathers progressively as people arrive. Dinner, informal conversation, shared dish if the table allows.

23:00 — Everyone goes home by their own means. No collective van, no management debrief, no obligatory extension. The signal sent is clear: the evening was yours, not the company's.

◆ THE BUDGET — WHAT IT ACTUALLY COSTS

For a team of 10 people:
— Hammam rented entirely (1h): €150 to €200 · *i.e. €15 to €20 per person*
— Walk: €0
— Belleville dinner (Asian, without excess): €15 to €25 per person
— **Total: €300 to €450 for 10 people — €30 to €45 per person**
Comparison: corporate two-day seminar for 10 people = €8,000 to €15,000. The evening costs 30 to 50 times less. It produces a neurochemical sequence the seminar cannot reproduce.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE NEUROCHEMICAL SEQUENCE — THE FOUR PHASES

WHAT HAPPENS IN THE BODY AND IN THE GROUP

The genius of this programme is that it chains four neurochemically distinct phases, each documented by psychology and neuroscience research, each preparing the next. This is not just an evening — it is an unwitting sequence of recovery and bonding.

◆ PHASE 1 — THE HAMMAM: CORTISOL DROP, OXYTOCIN RISE

The hammam's heat activates the parasympathetic system — the nervous system's "rest mode". Cortisol drops. The University of Granada is currently conducting a clinical trial specifically on the hammam experience's effect on biophysiological stress markers. The therapeutic touch of the scrub — well documented in massage therapy research — triggers oxytocin release and reduces ACTH, the hormone that stimulates cortisol production. Rapaport et al. (2012) established a dose-response relationship: repeated massage exposure durably lowers cortisol and vasopressin. One hour of shared hammam among colleagues produces the same social safety signal documented in research on loneliness and bonding: the body signals to the organism that it is not alone.

◆ PHASE 2 — THE FREE WALK: ENDORPHINS AND DETACHMENT

The 45-minute to one-hour walk from Barbès to Belleville is not a commute — it is a transition. Walking produces endorphins and lowers residual cortisol. The freedom of route — everyone chooses their path, companion, pace — sends an autonomy signal that enhances recovery quality. Tomasello and Wolf (Max Planck / Duke, 2023) document that walking together is one of humanity's unique social bonding mechanisms: sharing triadic experience in a common physical space creates a sense of closeness that prepares future cooperation. Sub-groups forming spontaneously during the walk are based on real affinities, not the org chart — precisely the informal bonds that reinforce operational cohesion.

◆ PHASE 3 — DINNER: SOCIAL BONDING OXYTOCIN AND SHARED MEMORY

The shared meal is documented by neuroscience research as a privileged context for human social bonding. Conversations around meals exchange primarily information about people — activating the "social brain" circuits more deeply than professional meetings. Oxytocin produced by positive interactions during dinner reduces cortisol and reinforces trust — exactly the relational capital the team will spend during the next production incident at 3am.

◆ PHASE 4 — AUTONOMOUS RETURN: THE SIGNAL OF RESPECT

Everyone goes home independently, with no obligation to extend. This final freedom is neurochemically important: it avoids the fatigue of a forced evening that transforms a positive memory into a constraint. It tells each team member that their personal time is respected — a psychological safety signal documented by Edmondson as a condition of collective high performance.

Amine RAITI · CC BY-NC-SA 4.0

NOT THE LUXURY SPA — THE NEIGHBOURHOOD HAMMAM

The choice of the neighbourhood hammam — not the hotel spa at €150 per entry — is a philosophical as much as an economic choice. It says something about the manager who organises it and about what they seek to produce.

◆ WHAT THE NEIGHBOURHOOD HAMMAM DOES THAT THE CORPORATE SPA DOES NOT

The corporate hotel spa is within the world of work. It has a reception desk with a badge, a laminated menu of services, towels embroidered with the logo. It subtly maintains the professional register — the environment says "you are on a business trip". The neighbourhood hammam in Barbès is in the city. It belongs to popular culture, to the neighbourhood, to a century-old tradition. The environment says "you are no longer at work." This contextual rupture signal is documented in psychology as a necessary condition for real recovery — the brain only disconnects from vigilance mode when the environment explicitly signals that this mode is no longer required.

◆ THE CULTURAL DIMENSION — THE HAMMAM AS MAÂLLEM OF CARE

The Moroccan or Turkish hammam is a social institution as old as the medina. It is not designed for performance or productivity — it is designed for collective recovery, socialisation and care. Sharing a hammam with colleagues is inviting the team into a culture of care that predates the world of work by several centuries. This temporal and cultural displacement — like the shared dish of "The Yemeni Dinner" — produces a level of dehierarchisation that corporate formats cannot reach.

◆ RENTING IT ENTIRELY — THE DETAIL THAT CHANGES EVERYTHING

Renting the entire hammam for one hour — €150 to €200 for 10 people — creates a private space that reinforces group safety. No outside gaze. No presence of strangers maintaining social vigilance. The team is among itself, in a defined space, with a defined duration. These conditions are exactly those researchers identify as favourable to oxytocin production and bond formation: safety, familiarity, absence of competition with strangers.

◆ NASSIHA — NOT A FORMAT FOR ALL CONTEXTS

The hammam presupposes a degree of comfort with physical proximity and the body. In a team that does not yet know each other well, or in cultural contexts where partial nudity is a sensitive subject, this format may not be the right first choice. The rule is the same as for Nerf: voluntary participation, without implicit pressure. The hammam is a format for a team that already has a minimal level of trust — it deepens this bond, it does not create it from scratch.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE FREE WALK — THE INVISIBLE LINK

BETWEEN THE HAMMAM AND THE DINNER — THE TRANSITION THAT MAKES THE BOND

The walk from Barbès to Belleville is the link that nobody plans and that is nonetheless the most precious part of the evening. It is not in the official programme — it is in the space between the hammam and the dinner. And it is precisely in this unstructured space that the strongest bond is built.

◆ WHAT THE FREE WALK PRODUCES THAT THE GUIDED WALK DOES NOT

A guided walk — "meet at 8:15pm outside Barbès metro, we walk together to Belleville" — would recreate a group structure that maintains the usual hierarchical and social dynamics. The freedom of route — everyone leaves when they want, with whom they want, by the path they choose — produces something different: the spontaneous formation of natural affinity sub-groups. The two colleagues who walk together because they live in the same direction, the three who stop for a coffee en route, the one who prefers to walk alone and arrives last — all these choices are respected. And this freedom produces conversations that group structure would have prevented.

◆ THE GEOGRAPHIC TRANSITION AS DISCONNECTION

Moving physically from Barbès to Belleville means crossing a part of Paris that belongs neither to business districts nor to conference zones and corporate restaurants. This route through popular, lively, multicultural Paris is in itself a disconnection from the professional environment. Environmental psychology documents that changing physical environment facilitates changing mental state. Walking through Paris is not a commute — it is open-air therapy.

◆ THE MEETING POINT COMMUNICATED IN ADVANCE — THE ORGANISATIONAL DETAIL

Communicating the restaurant address in advance — by message, not during a collective brief outside the hammam — creates real freedom of route. Everyone knows where to go. Everyone decides how to get there and with whom. Organisation is minimal, freedom is maximal. This "free meeting point" format is the same as that of hashers — runners who set an arrival point and freely choose their own path. The freedom of the route is the condition for the quality of the arrival.

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · WHAT THESE FORMATS DO NOT DO — AND WHAT THEY DO

PHYSICAL COHESION IS NOT FORCED COHESION

The formats documented in this corpus — Yemeni dinner, office Nerf, hammam + walk + dinner — share a characteristic that fundamentally distinguishes them from corporate team building: they do not claim to build cohesion. They create the neurochemical and social conditions in which cohesion builds itself.

◆ SOCIAL BUFFERING THEORY

Research on "social buffering" — documented in primates, rodents and humans — establishes that the presence of a familiar social partner reduces physiological responses to stressors. This mechanism is proposed as one of the evolutionary reasons for the formation of social relationships. Translated into SRE team terms: an engineer who has shared a hammam and dinner with colleagues manages the production incident physiologically differently from one who only knows colleagues from meetings. Physical familiarity reduces incident cortisol — not because it changes the difficulty of the problem, but because it changes the perception of the safety within which the problem is solved.

◆ OXYTOCIN AS OPERATIONAL TRUST CAPITAL

Positive face-to-face interactions produce oxytocin that reduces cortisol and reinforces trust. This trust capital is not stored in an HR database — it is stored in the body, in the neurological circuits of social memory. A team that has accumulated shared evenings has accumulated an oxytocin capital that is spent during moments of tension and collective decision-making. The person whose relaxed face you know from the hammam is the same person whose judgement you trust at 3am during an incident.

◆ WHAT THESE FORMATS DO NOT CLAIM TO DO

They do not resolve structural conflicts. They do not compensate for deficient management. They do not replace decent working conditions. As with Nerf, they are multipliers — they amplify what is already in place. A team with structural problems that shares a hammam will be a team with structural problems and hammam memories. Body care is a complement to management, not its substitute.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

THE BODY FIRST — EXTENDED QUARTERLY PROTOCOL

As a complement to the quarterly dinner documented in "The Yemeni Dinner", this study proposes adding an annual deeper format — hammam + walk + dinner — that produces a complete neurochemical sequence impossible to reach with a single format. The two formats are complementary and non-substitutable.

◆ **PROPOSED ANNUAL CALENDAR**

Quarter 1: Simple team dinner — thoughtfully chosen restaurant, bowling or walk beforehand. Light format, €78/person. Objective: bond maintenance.

Quarter 2: Hammam + free walk + neighbourhood dinner — the complete format of this study. €35/person. Objective: deep recovery and physical bonding.

Quarter 3: Competitive activity — karting, bowling, escape room, concert. Energetic format, fiero, adrenaline. Objective: activation and shared memory.

Quarter 4: Year-end dinner — longer, more generous. Informal annual review. Objective: closure and recognition.

Total annual budget per person: €200 to €300. The cost of a single day of corporate seminar per person — for four evenings covering the full neurochemical spectrum of cohesion.

◆ **OTHER FORMATS IN THE SAME SPIRIT — BEYOND THE HAMMAM**

The hammam is one format among others in the same "body first" logic:
— **Rented pool** (one hour outside public hours) → same decompression and dehierarchisation effects, summer format
— **Group massage session** (practitioner comes to the premises) → less travel, effective for teams with logistical constraints
— **Opera or concert** → collective shared emotion, fiero of the grand hall, memorable shared experience
— **Karting or paintball** → adrenaline, fiero, light competition — same logic as Nerf but in event format
The rule common to all these formats: the body lives something together before the table brings minds together.

HUMAN

◆◆◆

€35 per person. A hammam rented entirely in Barbès. A free walk to Belleville. An Asian dinner. Everyone goes home fresh and recharged. The neurochemistry is the same as that of €2,000 seminars — the budget is not.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE PAGER AND THE BODY

What permanent on-call destroys in silence

HUMAN

◆ **CONTEXT**

This study extends "The Economy of Silence" (uptime as the sole SRE metric) and "The Departure of the Last One Who Knows" (mechanism 4 — burnout by accumulation). It addresses what permanent on-call does to the body and cognition over time — a reality documented by occupational medicine, ignored by the infrastructure industry. This document is not a complaint. It is a risk analysis applied to the human resource.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · WHAT ON-CALL DOES TO THE BODY

THE 3AM WAKE-UP AND ITS PHYSIOLOGICAL CONSEQUENCES

The pager sounds at 3am. The engineer wakes abruptly, cortisol spiking, heart rate at 100. They open their terminal, analyse alerts, make decisions on a production system. This scenario is normal in infrastructure. It is documented by occupational medicine as one of the most aggressive known disruptors of the circadian cycle — on a par with rotating night shifts, hospital duty, and shift work.

◆ NOCTURNAL COGNITIVE DEGRADATION — WHAT SCIENCE SAYS

Occupational medicine and neuroscience research converges: cognitive capacities between 2am and 6am are significantly reduced — extended reaction time, diminished complex reasoning, reduced working memory, tendency to fixate on incorrect solutions. An incident managed at 3am by an engineer woken abruptly is statistically more likely to produce a human error, a suboptimal decision, or an unnecessary escalation than an identical incident managed during the day.

◆ WHAT THE SECTOR REFUSES TO MEASURE

The infrastructure industry measures MTTR — Mean Time to Repair. It does not measure the correlation between incident trigger time and resolution quality. It does not measure the number of incidents that escalate because a 3am decision worsened the initial problem. It does not measure the cost of nocturnal human errors. This cost is real. It is simply invisible in operational performance dashboards.

◆ NASSIHA — THIS IS NOT A CRITIQUE OF ON-CALL

On-call is an operational necessity in critical production environments. This study does not call for its elimination. It calls for its recognition as an operational risk, and its management as such — with the same analysis and mitigation tools one would apply to any other infrastructure risk.

HUMAN

2

SECTION 2 · THE SILENT ACCUMULATION

BURNOUT IS NOT AN EVENT — IT IS AN ACCUMULATION

Burnout in infrastructure does not look like a sudden breaking point. It looks like a progressive, invisible erosion that neither the engineer nor their manager sees coming, because it occurs below the detection threshold of normal indicators. The engineer continues responding to alerts, managing incidents, delivering. They do so with increasingly reduced cognitive and emotional resources — but they do it, and this continued capacity masks the ongoing degradation.

◆ THE THREE DOCUMENTED STAGES

Stage 1 — Active compensation: the engineer compensates for fatigue through conscious effort. They are more careful, double-check, document better to offset the memory gaps they are starting to feel. This stage is invisible externally — they will overperform on indicators.

Stage 2 — Progressive disengagement: compensation becomes too costly. The engineer begins doing the bare minimum — not out of laziness, but to preserve remaining resources. Work quality degrades slightly. They begin avoiding complex tasks.

Stage 3 — Collapse: accumulation reaches a threshold. The engineer stops — sick leave, resignation, or the most dangerous scenario for the organisation — they stay but are no longer really there.

◆ ON-CALL AS AN AMPLIFIER

On-call is not the sole cause of burnout in infrastructure. But it is a powerful amplifier of all other causes — excessive workload, lack of recognition, loss of meaning, organisational conflicts. An already pressured engineer who accumulates on-call nights accelerates their progression to stage 3. On-call frequency is as reliable a burnout risk indicator as any workplace wellbeing questionnaire.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · MISSING METRICS

WHAT THE SECTOR DOES NOT MEASURE

The infrastructure industry has developed sophisticated metrics to measure the health of its systems — SLA, MTTR, MTTD, error rate, latency. It has not developed equivalent metrics to measure the health of the engineers who maintain these systems. This asymmetry is revealing: the organisation knows how to measure what costs it directly (downtime, incidents) but not what will cost it tomorrow (team erosion).

◆ THE METRICS THAT SHOULD EXIST

— On-call frequency per engineer per month (number of nights and weekends)
— Number of nocturnal alerts received per engineer (not just incidents handled)
— Average recovery time after an on-call night (hours before returning to normal work)
— Rotation rate on on-call positions (high rate is an early burnout signal)
— Correlation between night-managed incidents and human errors documented in post-mortems

None of these metrics are standard in the sector. Some exist in the most mature organisations. The majority do not calculate them.

◆ THE LINK WITH "THE ECONOMY OF SILENCE"

"The Economy of Silence" documented that uptime is the only metric that counts in many infrastructure organisations — that anything not visible in the availability dashboard is invisible. Engineer erosion is the textbook case of this invisibility: it generates no alert in monitoring tools. It generates a resignation six months later, or a major incident whose true root cause nobody will understand.

◆ NASSIHA — METRICS ALONE ARE NOT ENOUGH

Measuring on-call frequency is necessary but not sufficient. An organisation that measures and changes nothing produces data without value. Measurement only makes sense when coupled with pre-defined action thresholds — beyond X on-call nights per month, rotation is mandatory; beyond Y nocturnal alerts per week, the alerting system is reviewed.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

OPERATIONAL HEROISM AS A TOXIC NORM

The infrastructure sector has developed a culture of operational heroism — the implicit valorisation of the engineer who manages the 3am incident without complaining, who is always present, who "holds" whatever happens. This culture is understandable in its origin: critical systems require reliable people. It is toxic in its effects: it implicitly punishes those who set limits and rewards those who wear themselves out without saying so.

◆ HOW HEROISM MASKS STRUCTURAL PROBLEMS

A heroic engineer absorbing an excessive on-call load makes the problem invisible to their manager. The system runs. Incidents are resolved. Metrics are good. The underlying structure — too few engineers, overly sensitive alerting, fragile architecture — is never questioned because individual heroism compensates for its flaws. Until the engineer stops.

◆ THE BLAMELESS POST-MORTEM — AND ITS LIMITS

SRE culture introduced the blameless post-mortem — incident analysis without designating an individual culprit, seeking systemic causes. This is a real advance. But the blameless post-mortem addresses technical incidents. It does not address human incidents — burnout, the departure of an expert, the error of an exhausted engineer. These events are rarely treated with the same analytical rigour as server faults.

◆ WHAT MATURE SRE CULTURE DOES DIFFERENTLY

The most mature SRE organisations — certain teams at Google, Netflix, Stripe — have begun treating on-call load as an operational risk on a par with technical debt. They set maximum on-call thresholds per engineer, make rotation mandatory, measure nocturnal alert volume and reduce it as they would pay down technical debt. This is not goodwill — it is risk management.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

TREATING ON-CALL AS AN OPERATIONAL RISK

The levers that follow are not wellbeing measures — they are operational risk management measures. The distinction matters: it shifts the conversation from goodwill to reliability. An exhausted engineer is an operational risk on a par with a server without redundancy.

◆ LEVER 1 — MAXIMUM ON-CALL THRESHOLD PER ENGINEER

Define and respect a maximum number of on-call nights per engineer per month. This threshold is not a preference — it is an operational limit, like a server's maximum load limit. If the team cannot respect this threshold, it is an understaffing signal that should trigger recruitment or a reduction of the on-call scope.

◆ LEVER 2 — REDUCE NOCTURNAL ALERT VOLUME

The majority of nocturnal alerts are either false positives or alerts for events that do not require immediate intervention. Reducing this volume is engineering work — revising alert thresholds, removing alerts that never trigger action, consolidating redundant alerts. This work is systematically deprioritised. It should be treated as technical debt to be actively paid down.

◆ LEVER 3 — REAL FINANCIAL AND TIME COMPENSATION

On-call has a real physical and cognitive cost. It must be financially compensated at a level that reflects this cost — not symbolically. It must also entitle the engineer to protected recovery time the following morning: an engineer woken at 3am cannot be in a meeting at 9am without performance degradation. This recovery is not a luxury — it is the recharging of a critical operational resource.

◆ LEVER 4 — ARCHITECTURE AS AN ON-CALL REDUCER

The best response to excessive on-call is not organisational — it is architectural. Self-healing systems, resilient architectures, well-designed automatic failovers reduce the number of incidents requiring nocturnal human intervention. Investing in architectural resilience is the most direct investment in the health of on-call teams.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

RECOGNISING THE ENGINEER AS CRITICAL INFRASTRUCTURE

The central proposal of this study is simple: the on-call engineer is a component of the infrastructure as critical as the servers they maintain. They deserve the same level of monitoring, redundancy and preventive maintenance as those servers. A server without redundancy is an unacceptable risk. An on-call team without rotation is the same risk applied to the human resource.

◆ MEASURE 1 — PUBLISH ON-CALL METRICS IN SOCIAL REPORTS

Companies with more than 50 employees produce an annual social report. Integrating on-call metrics into it — average number of nights per engineer, nocturnal alert volume, rotation rate on on-call positions — makes visible what is structurally invisible. This transparency creates compliance pressure that the absence of measurement does not allow.

◆ MEASURE 2 — ADDRESS ON-CALL IN COLLECTIVE AGREEMENTS

The Syntec and digital sector collective agreements frame overtime, flat-rate day arrangements, and the right to disconnect. They do not specifically frame recurrent nocturnal on-call in critical production environments. A specific clause — maximum threshold, minimum compensation, guaranteed recovery time — in digital sector agreements is a measure within reach of social negotiation.

◆ MEASURE 3 — INTEGRATE ON-CALL LOAD INTO MANAGEMENT EVALUATION CRITERIA

A manager evaluated on system availability but not on their team's on-call load has an incentive to understaff teams and overload individuals. Adding on-call load as a management evaluation criterion — at the same level as technical indicators — changes the structural incentives of middle management.

◆◆◆

The pager does not sound in a vacuum. It sounds in a body. That body has limits. Ignoring them is an operational risk — not proof of resilience.

◆◆◆

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE BODY THAT MAKES

From LED to Matter

Arduino · Laser · CNC · 3D Printer in the Educational Arsenal

◆ POSITION OF THIS STUDY

This study is the direct continuation of "Deamputation at the Source". The Arduino in secondary school was the first contact — code that controls an electrical signal. This study proposes the second level: code that creates matter. Laser, CNC, 3D printer — three machines, one continuous thread with the Arduino. And an annual project: building the machine itself, to understand what it is before programming it.

MACHINES	AGE	BUDGET	WATERMARK
3	11-18 YEARS	€1,080 OF €2,400	HUMAN

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

HUMAN

I

SECTION I · THE COMPLETE CAUSAL CHAIN

FROM LED TO MATTER — ONE CONTINUOUS THREAD

"Deamputation at the Source" established the Arduino as first contact — code that controls an electrical signal, an LED that blinks, a sensor that reads temperature. That is level 1. This study proposes level 2: the same Arduino, the same code, but instead of controlling an LED, it controls a motor. And that motor moves a tool. And that tool transforms matter.

◆ THE CAUSAL CHAIN — STEP BY STEP

Step 1 — Arduino + LED (Year 7): digitalWrite(13, HIGH) → current → LED on. Code controls electricity. The student understands their instructions do something in the physical world.

Step 2 — Arduino + stepper motor (Year 8): stepper.step(200) → pulses → 200 steps → precise 360° rotation. Code controls a mechanical displacement. The student understands electricity can serve movement.

Step 3 — Arduino + XYZ axis (Year 9): G-code X25 Y30 → motor X advances 25mm, motor Y advances 30mm. Code controls a position in space. The student understands mathematical coordinates exist in matter.

Step 4 — Arduino + tool (Year 10): the laser engraves, the router cuts, the extruder deposits filament. Code creates matter. The student is no longer programming — they are making.

◆ WHAT THE CHAIN SAYS THAT THE ARDUINO ALONE DOES NOT

The Arduino alone can seem abstract. The chain answers the question students always ask: "What is this for?" The Arduino controlling a CNC that engraves the student's name on wood — that question disappears. Code has produced a real object. This is the corpus founding thesis: *the metal precedes the code*. These machines are its most physical demonstration.

◆ NASSIHA — THE ARDUINO IS NOT REPLACED, IT IS PROMOTED

The CNC3018, the 3D printer and the laser do not replace the Arduino — they use it. The brain of all three machines is an Arduino or derivative board (GRBL for the CNC, Marlin for the 3D printer). The student who learned Arduino in year 7 recognises its architecture in year 10 when they open the CNC control board. The continuity is total.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE THREE MACHINES — WHAT EACH ONE BRINGS

THREE TOOLS · THREE DOORS · ONE G-CODE

MACHINE	WHAT IT DOES	PRICE	LEVEL
2D Laser engraver	Engraves wood, leather, cardboard, slate. Immediate, beautiful, personal result. Ideal entry door.	€100-150	Year 8 / Year 10
CNC3018 (3 modes)	Laser + CNC router + drawing pen. Same G-code, three results. The pedagogical machine par excellence.	€150-200	Year 9 / Sixth form
3D Printer (Ender-3)	Creates matter ex nihilo. Volume, gravity, tolerances, supports. The deepest understanding of the code-object link.	€200-300	Year 11 / Year 12

◆ THE CNC3018 — THE PEDAGOGICAL MACHINE PAR EXCELLENCE

The CNC3018 is the richest pedagogically because it has three modes on the same frame — laser, CNC router, drawing pen. Same G-code, three different physical results. The laser burns the surface. The router cuts into wood or aluminium. The pen traces on paper. The student who runs all three modes with the same file understands something fundamental: code is neutral — it is the tool that decides the physical result. This is the first lesson in concrete abstraction.

◆ THE SCIENTIFIC BASIS — WHAT PSYCHOLOGY SAYS ABOUT THE FEELING OF MAKING

This feeling of satisfaction is not subjective — it is scientifically documented through three distinct phenomena.

The IKEA Effect (Norton, Mochon & Ariely, 2012 — Journal of Consumer Psychology): Four studies on subjects assembling IKEA boxes, folding origami, and building Lego sets demonstrate that individuals value their personal creations as highly as expert creations — even when objectively less well-made. The effort of making produces emotional attachment to the created object. This effect only applies to tasks completed to the end: if the creation is destroyed before completion, the satisfaction disappears. This is why the project must reach the finished object — the engraved part, the completed print, the machine's first movement.

The Self-Creation Effect (Brunneder & Dholakia, 2018 — Marketing Letters): Seven field and laboratory studies show that when a person self-creates a product, they appreciate it more, consume it more mindfully, and experience greater domain-specific and general well-being. This effect is amplified by self-consciousness — the student who knows they made something with their own hands develops a self-esteem that passive consumption does not produce.

The Maker Movement and Subjective Well-Being (Journal of Happiness Studies, 2017): A study of 465 students shows that a "maker" identity — perceiving oneself as someone who makes things — is a significant predictor of subjective well-being. Making activities (sewing, cooking, electronics, DIY) produce a sense of accomplishment and competence that transfers to other life domains.

What this means for the student: when the 13-year-old girl sees her name appear on a piece of wood engraved by the laser she programmed, she does not simply feel pride — she feels a neurological attachment to her own competence. This feeling, once anchored, does not disappear. It becomes the basis of a lasting curiosity about the physical world.

◆ THE LASER — THE DOOR FOR GIRLS

The laser engraver is the most effective entry door for profiles who do not yet project themselves into engineering. Because the result is beautiful, immediate, and personal. Engraving a name on wood, creating a leather jewel cut by laser, personalising a notebook — these are projects that attract profiles the Arduino alone does not attract. And behind these projects: the same G-code, the same XYZ coordinate system, the same fundamental competence.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE XYZ COORDINATE SYSTEM — THE BRIDGE BETWEEN MATHS AND MATTER

THREE LETTERS ON THE BOARD · THREE AXES IN THE MACHINE · ONE TRUTH

The maths teacher has been teaching the XYZ Cartesian coordinate system for years. Students learn to place points in space, calculate distances, find intersections. It is abstract — crosses on squared paper. The day the student types G-code X25 Y30 Z-2 into the CNC terminal and sees the router go to exactly that position, something triggers. The morning's coordinates exist in matter that afternoon.

◆ WHAT EACH AXIS TEACHES

X AXIS — WIDTH: horizontal displacement left to right. In maths: the abscissa. In the CNC: the X motor advances the head 25mm to the right. The student sees the abscissa move physically.

Y AXIS — DEPTH: horizontal displacement front to back. In maths: the ordinate. In the CNC: the Y motor positions the bed. The ordinate has mass, inertia, mechanical noise.

Z AXIS — HEIGHT: vertical displacement. In maths: the Z coordinate. In the CNC: Z-2 means the router descends 2mm into the material. The coordinate has a depth of cut, a cutting force, a risk of breakage if miscalculated. The coordinate becomes an engineering decision.

◆ THE DISCIPLINE BRIDGE — MATHS + PHYSICS + TECHNOLOGY

The maths teacher teaches the XYZ coordinate system, coordinates, displacement vectors. On paper, without any machine.

The physics teacher explains the stepper motor — current, magnetic field, torque, angular step. Why the motor stops exactly where told. The physics of positioning.

The technology teacher runs the machine — G-code, control software (GRBL), feed rate, depth of cut. The interface between code and matter.

Together, the three teachers around the same machine give the student what none of the three can give alone: integrated understanding of the link between mathematical abstraction, physical reality, and technical decision.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE LIVING PROOF — THE RSCNC32

THE M'ALLEM WHO SHOWS THEIR MASTERPIECE BEFORE ASKING FOR YOURS

The corpus does not speak in a vacuum. The RSCNC32 is a CNC built by hand, from scratch, by Amine RAITI. The mechanical parts were manufactured with that same machine. The photos of these parts are the physical demonstration that everything this study proposes is not only possible — it has already been done. Without industrial budget. Without a team. With hands, time, and the decision to understand rather than to purchase.

① Frame under assembly

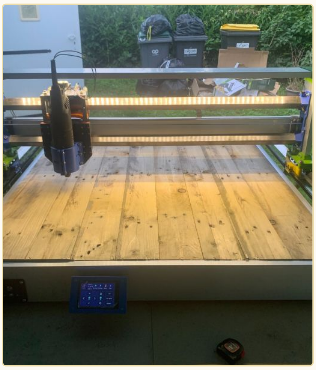

② Control board, drivers and wiring

③ Engraved parts produced with the machine

④ Moroccan Gibs decoration — engraved on the house

◆ RSCNC32 — HAND-BUILT · AMINE RAITI

CNC machine built from scratch — frame, wiring, GRBL firmware, calibration, all done by hand. The Moroccan Gibs-style decorations were engraved with this same machine on panels of the house. Proof that autonomous fabrication is possible, accessible, and already done — without industrial budget, without a team.

◆ WHAT THE HOMEMADE RSCNC32 SAYS THAT CATALOGUES DO NOT

A hardware catalogue says: here is a CNC for €200. It does not say how it works. It does not say what it means to choose between a Nema 17 and a Nema 23 motor. It does not say why the mechanical rigidity of the frame determines the precision of the result. It does not say what one learns by wiring motor drivers oneself, calculating the power supply current, installing the GRBL firmware.

The homemade RSCNC32 says all of this — not with words, with engraved parts. The student who sees these parts and learns they were manufactured with a hand-built machine understands that autonomous fabrication is possible. Not reserved for diploma-holding engineers. The hands that know precede the title that says so.

◆ THE M'ALLEM AND THEIR MASTERPIECE — THE COMPAGNON LOGIC

In the Compagnons du Devoir tradition, the master shows their masterpiece before asking for the aspirant's. Not vanity — pedagogy. Proof that the path has been walked, that the result is real, that transmission is possible. The RSCNC32 is that masterpiece in the context of this study. It says to the student: here is what your hands can produce. Not in twenty years after a degree. In one school year, with a pooled teacher budget.

◆ NASSIHA — THE BUILT MACHINE UNDERSTANDS THE BOUGHT MACHINE BETTER

The student who has built a CNC — even simple, even imperfect — understands why the purchased CNC has those dimensions, those tolerances, those limits. They do not use it as a black box. They use it as someone who knows what is inside. This is the difference between the operator and the maker. Between the DevOps who uses the cloud and the SRE who knows what is under the cloud.

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · SECONDARY AND SIXTH FORM — PROGRESSION BY LEVEL

FIVE YEARS · FIVE LEVELS · ONE COMPLETE CHAIN

◆ YEAR 7 — ARDUINO (LED, SENSOR)

First contact. Code → electrical signal. The LED as revelation. Deamputation begins here. Each student alone with the kit.

◆ YEAR 8 — LASER ENGRAVER

Name tag project (45 min). Student draws in Inkscape, generates G-code, engraves on wood. Object to take home. Door for girls: leather jewel, personalisation.

◆ YEAR 9 — CNC3018 (3 MODES)

Same G-code, three results. The XYZ coordinate system becomes real. Router, laser, pen. The student understands code is neutral — the tool decides.

◆ YEAR 10 — 3D PRINTER

Creation ex nihilo. FreeCAD modelling → Cura → Ender-3. Volume, gravity, supports, tolerances. The deepest understanding of the code-object link.

◆ THE ANNUAL PROJECT — YEARS 11-12 — BUILDING THE MACHINE

The most ambitious and most formative level. A group of 4 to 6 students builds an open source CNC or 3D printer over the school year. What this project teaches:

Stepper motors: current, torque, angular step, driver. Understanding that 200 steps = 1 turn = precise displacement of Xmm according to the lead screw pitch. The physics of positioning.

The control board (GRBL / Marlin): Arduino firmware that receives G-code, decodes it, and sends pulses to the drivers. Students recognise the Arduino from year 7 — it is now the brain of a machine.

Power supply: each motor consumes. Each driver heats. Each power supply choice carries responsibility. Miscalculating the power supply breaks hardware. The physics of power becomes an engineering decision.

Mechanics: machining precision, mechanical play, frame rigidity. A misaligned axis produces a badly engraved part. Code does not compensate for faulty mechanics.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE ANNUAL PROJECT — WHAT THE STUDENT UNDERSTANDS THAT NOBODY ELSE DOES

BUILDING THE MACHINE — THE FIRST MASTERPIECE

There is a fundamental difference between the student who has used a CNC and the student who has built one. The one who built it knows what is inside. They do not use a black box — they use something they understand at every level. This difference, twenty years later, is the difference between the engineer who configures a cloud service and the engineer who understands what the cloud service hides.

◆ ANNUAL PROJECT SCHEDULE — CONCRETE MILESTONES

September-October — Architecture: choose the open source model (MPCNC, Voron, Prusa i3). Read the documentation. Make the parts list. Calculate the budget. Divide roles — mechanical, electronics, firmware, testing.

November-December — Procurement and mechanics: order parts (AliExpress, local suppliers). Print missing parts on the school's printer. Assemble the frame. Understand why frame rigidity determines final precision.

January-February — Electronics: wire the drivers, motors, end stops, power supply. Identify wiring errors (short circuits, wrong polarity). Understand why a poorly cooled driver overheats and cuts out.

March-April — Firmware: flash GRBL or Marlin. Configure parameters — steps/mm, motor current, maximum speed, acceleration. Run first homing and first manual moves.

May-June — Calibration and masterpiece: engrave the first real part. Correct it. Redo it. Present the machine to other students and teachers. The machine stays in the school — it will be used by students in future years.

◆ THE MACHINE STAYS — TRANSMISSION BEGINS

The machine built by year-12 students stays in the school. Year-10 students will see it, use it, perhaps improve it. Year-11 students will maintain it. This is compagnon transmission applied to the state school. The masterpiece does not disappear with the cohort — it becomes the tool of the next one.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · POOLED BUDGET LEVEL 2 — COMPLETE INVENTORY

€2,400/YEAR · €1,080 FOR MACHINES · €1,320 RESERVE

◆ INVENTORY — WHAT €2,400/YEAR BUYS (MACHINES + ARDUINO)

20 Arduino Starter Kits (documented in "Deamputation at the Source"): €700 — class of 20 students, hardware reusable 5 years.

5W laser engraver (xTool D1 or equivalent): €150 — engraves wood, leather, cardboard. Immediate, beautiful, personal result.

CNC3018 Pro (3 modes: laser + router + pen): €180 — the pedagogical machine par excellence. Three modes, one G-code.

3D Printer Ender-3 V2: €250 — creation ex nihilo. Robust, documented, huge support community.

Open source CNC build kit (MPCNC or equivalent): €200 — Nema 17 motors, A4988 drivers, Arduino UNO board, aluminium profiles, hardware, belts.

Year 1 consumables: €300 — plywood, PLA filament, slate plates, leather, wiring, breadboards.

Software: €0 — Inkscape (vector), FreeCAD (3D), Cura (slicer), GRBL (CNC firmware), LaserGRBL (laser control). All open source, all free.

Total machines: €1,080

Year 1 reserve: €620 (additional consumables, spare parts, contingencies)

Year 2 reserve: the full €2,400 of year 2 goes entirely to consumables and improvements — the machines are already there.

◆ WHY THIS BUDGET — AND NOT A GRANT APPLICATION

This pooled budget proposal is deliberately independent of any administrative process. No project call. No academic dossier. No application to a local council, regional authority or ministry. These channels exist — and they can complement this arrangement if they succeed. But they take time that children do not have: the bifurcation of year 7 does not wait for a ministerial decree.

The pooled budget says one simple thing: we are responsible for what happens in our classrooms. We are not waiting for anyone to act.

The parents' association — a natural complementary lever: parents have a direct and legitimate interest in their children's education. A parents' association that decides to fund digital fabrication equipment is not organising a school fair — it is investing in the technical future of its children. A voluntary contribution of €5 per family from 50 families adds €250 to the annual budget. From 100 families: €500. Without any additional administrative structure. The parents' association already exists in every school — it only requires teachers and parents to decide together that this priority is worth it.

◆ THE COMPLETE CHAIN — FROM LED TO SOVEREIGNTY

Year 7: Arduino + LED → code controls electricity.

Year 8: Laser engraver → code controls light, which engraves matter.

Year 9: CNC3018 → code controls XYZ axes, which shape matter.

Year 10: 3D Printer → code creates matter ex nihilo.

Years 11-12: Build the machine → the student understands what all machines hide.

The student who has traversed this complete chain understands, at 18, what most software engineers do not understand at 30: that code does not exist in a vacuum — it exists to transform matter. That technical sovereignty begins in the hands. That the metal precedes the code.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE MENTOR AND THE METAL

MAÂLLEM · METAÂLLEM · SRE

HUMAN

◆ CONTEXT OF THE STUDY

The Zellij of Fès, the Gibs of Meknès, the leather craft of Marrakech — these arts have survived only because the Maâllem/Metaâllem process was never broken. This process is the same one that produced the best infrastructure engineers this corpus has encountered — not in classrooms, but in server rooms. This study argues that the mentor is the primary channel for tacit knowledge transmission in infrastructure, and that without them, even the most rigorous training remains incomplete.

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE MAÂLLEM AND THE MEDINA

WHAT MOROCCO HAS UNDERSTOOD FOR CENTURIES

In the Moroccan medinas, the Maâllem — literally "the one who knows" — is the master craftsman who holds knowledge acquired over decades of practice. He has been laying Zellij tiles for forty years. He knows how the hammer must sound against the faience for the cut to be clean. He knows which clay for which gibs depending on seasonal humidity. He knows how Marrakech leather responds to August heat.

This knowledge is not in a manual. It is not in a skills framework. It is in the hands, in the eyes, in the body of the Maâllem — and it can only be transmitted one way: the Metaâllem, the apprentice, sits beside the master. He watches. He imitates. He fails. He starts again. And one day — after months, sometimes years — the gesture becomes right. Not because it was explained. Because he watched and repeated until it was incorporated.

◆ WHY THESE INDUSTRIES SURVIVED

Morocco's ancestral arts survived 20th-century industrialisation, globalisation and mass tourism — not despite their refusal to modernise, but through the continuity of the Maâllem/Metaâllem process. When this chain broke in certain guilds — through rural exodus, forced industrialisation, absence of transmission — the art disappeared. Not the techniques in the books. The art. Because art is not in books. It is in the hands that learned it from other hands.

The Maâllem at work — from the laying hand to the completed masterpiece · Moroccan medina

◆ THE SAME MECHANISM — ANOTHER MEDINA

The server room is a medina. The cables are the alleyways. The racks are the workshops. The senior SRE who shows the junior how to listen to fan noise to detect a thermal anomaly before the alert sounds is a Maâllem. He does not explain it. He does it. And the junior who watches, who asks questions, who repeats under supervision, becomes in turn a Maâllem — capable of transmitting.

Amine RAITI · CC BY-NC-SA 4.0

WHAT THE SENIORS TRANSMITTED THAT THE COURSES DID NOT

When I joined Ecritel, I was not a beginner. I had the training, the diplomas, the foundations. But training had not taught me what the seniors in place transmitted in real production conditions. What I received from them appears in no skills framework. Yet it is what built the engineer I became.

◆ THE THREE TRANSMISSIONS — FROM ECRITEL TO WEBORAMA

Patience. A senior taught me not to act immediately on an incident. To observe first. To read what the system was saying before touching anything. This discipline — not panicking, not fixing too fast, understanding before intervening — cannot be taught in a classroom. It is transmitted through the example of a man who stays calm when everyone around wants to act. I watched him do it. I integrated the posture before understanding the reason.

Technologies. Not technologies as in courses — technologies as they actually behave in production. The edge cases, the undocumented behaviours, the things that work when they should not and vice versa. This knowledge accumulates over years of incidents. A senior who shares it condenses years of experience into a few hours of informal conversation around a fault.

Management capabilities. How to write a post-mortem that does not seek a culprit. How to announce a failure to a client without losing their trust. How to manage an on-call team without creating resentment. These competences are acquired by observing managers who exercise them — not by reading management books.

◆ NASSIHA — WHAT THIS TRANSMISSION REQUIRES

The Maâllem/Metaâllem transmission requires time, proximity and trust. It cannot be decreed by an HR policy. It requires the senior to have the availability to transmit — which presupposes a work environment structured not solely around immediate performance, but also around the continuity of knowledge.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

TACIT KNOWLEDGE — THE INVISIBLE LAYER OF COMPETENCE

Training — whatever its quality — transmits explicit knowledge. What can be formalised, written, assessed. The Foundation of Iron transmits cabling, addressing, protocols, systems administration. A good course transmits the fundamentals of theory and supervised practice. What training cannot transmit is tacit knowledge — the invisible layer of competence that only real-condition practice, under a mentor's eye, can build.

◆ TACIT KNOWLEDGE IN INFRASTRUCTURE — WHAT IT IS CONCRETELY

— **Incident intuition:** the senior who enters the server room and knows, before looking at metrics, that something is wrong — because the ambient sound has changed, because a switch LED is blinking in a particular way, because the room temperature is half a degree higher than usual. This intuition is acquired over thousands of hours of physical presence in server rooms.
— **Judgement under pressure:** knowing when to act and when to wait. Knowing when to escalate and when to manage alone. Knowing when the quick fix will worsen the situation. This judgement is not taught — it is forged in real incidents, alongside someone who has already forged it.
— **Team culture:** how one behaves in an SRE team. Where the boundary lies between help and dependence. What the difference is between productive silence during an incident and anxious silence. The right way to ask a senior a question without signalling you have not understood what you should have. These cultural codes are learned through immersion.

◆ WHY THE COURSE ALONE IS NOT ENOUGH — EVEN THE BEST

A driving lesson can be excellent — highway code theory, simulator exercises, explanation of reflexes to develop. It does not replace the 3,000 kilometres with an experienced instructor in real traffic conditions. Bare-metal infrastructure works exactly the same way. The Foundation of Iron is the highway code. The production server room with a senior is real traffic. Both are necessary. Neither alone is sufficient.

HUMAN

◆ NASSIHA — THE MENTOR IS NOT A TRAINER

The mentor does not follow a programme. They do not validate competences. They have no formalised pedagogical objectives. They do their job — and the Metaâllem is there, alongside, watching and learning. This absence of formal structure is precisely what makes the transmission effective. The junior learns in real time what the senior actually does — not what a programme says they should do.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE OBSERVATION DAY — THE MINIMAL FORMAT THAT CHANGES EVERYTHING

ONE DAY IN A SERVER ROOM IS WORTH MORE THAN A SEMESTER OF THEORY

The SRE/student mentorship does not need to be a six-month internship. It can start with one day. Or two. A day where the student is present as an observer — not as an intern, not as productive, not as assessable. As an observer. They watch. They are there. And this presence alone changes something irreversible in their mental representation of the profession.

◆ WHAT HAPPENS DURING AN OBSERVATION DAY

Morning: the student arrives at the premises. They see a technical open space — the screens, the open terminals, the working atmosphere. They attend a stand-up. They hear the tickets, the priorities, the vocabulary. Not to understand everything — to hear the register, the way engineers talk about systems, the casualness or tension depending on the signals of the moment.

The rack visit: the SRE takes them into the server room. They see the servers. They hear the fans. They feel the temperature difference between the hot and cold aisles. They touch a cable. They watch an LED blink. This physical moment — this sensory encounter with the infrastructure — is often the first trigger. The system stops being an abstraction. It becomes real.

The break: the student eats with the team. The informal conversations are the most valuable part of the day. Stories of incidents, anecdotes about memorable failures, debates about architecture choices, corridor jokes. They hear the team culture in natural mode. They understand how these people think about their profession when nobody is evaluating them.

The incident if luck allows: if an incident occurs during the day, the student observes the response in real time. This is the most formative possible sequence — watching how an experienced engineer moves from alert to diagnosis, from diagnosis to decision, from decision to resolution. Without being in the critical path. Just observing.

◆ WHAT THE STUDENT TAKES AWAY — THAT TRAINING WOULD NOT HAVE GIVEN

A concrete image of the profession that replaces abstract or romanticised representations. The certainty that this profession is accessible — that the engineers in front of them are normal people with learned methods, not born geniuses. A human connection with a professional who can answer their questions. And often — the vocational trigger that definitively confirms or rules out the professional direction.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

FROM INFORMAL MAÂLLEM TO STRUCTURED MENTORSHIP

Maâllem/Metaâllem mentorship in infrastructure does not need to be a heavy programme. It needs to be intentional. The difference between the informal mentoring that already exists in teams and structured mentorship is the explicit decision to do it — and the minimal framework that makes it possible without overloading the mentor.

◆ THE ROLE OF HR — FACILITATOR, NOT MANAGER

HR is the bridge between the school or training organisation and the infrastructure team. Their role is not to manage the content of mentorship — they do not have the technical competence for that. Their role is to create the connection: identify SREs and engineers willing to mentor, contact BTS SIO programmes, DUT Networks, engineering schools and training organisations, and logistically organise the observation days. A simple agreement, a day blocked in the calendar, a visitor badge. The rest happens naturally.

◆ WHAT WE ASK OF THE SRE MENTOR — AND WHAT WE DO NOT

We ask: to do their work normally in the student's presence. To take them on a server room visit. To include them in the break. To be available to answer their questions at the end of the day. That is all.

We do not ask: to prepare a pedagogical programme. To validate competences. To fill in assessment forms. To adapt to a framework. To modify their working method. The power of the Maâllem format is precisely that knowledge transmits in normal action, not in an artificial pedagogical parenthesis.

◆ THE NATURAL PROGRESSION — FROM THE DAY TO APPRENTICESHIP

One observation day can become two. Then a week of observation placement at the end of the BTS. Then a work-study programme. Then a first permanent contract in the team the student already knows — because they have seen the racks, eaten with the engineers, observed the team culture. Recruitment is no longer a gamble on a CV — it is the logical continuation of a built relationship. And integration is divided by ten because the Metaâllem already knows their Maâllem.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL — THE COMPLETE CHAIN

TRAINING · MENTOR · MENTORSHIP · EMPLOYMENT

This study completes the Opération Dindon corpus by adding the missing link between training and employment. The Foundation of Iron transmits the fundamentals. The BTS and DUT provide the academic framework. Arduino in middle school opens vocations. But between the best-designed training and the operational engineer, the Maâllem is missing — the transmitter of tacit knowledge that only real production conditions can activate.

◆ THE IDEAL CHAIN — FROM 13 TO OPERATIONAL

Middle school (13-15): Arduino, robotics, connected car. Hands in the machine before theory. The vocational trigger.

Vocational high school / BTS / DUT (16-20): The Foundation of Iron in the frameworks. Fundamentals from physical layer to service layer. Basic technical competence.

Observation day (during BTS / DUT): One day in a real server room. The racks, the teams, the break, the incident if luck allows. The concrete that anchors training in reality.

Work-study or long placement (end of training): The Metaâllem in real conditions, alongside a Maâllem. Tacit knowledge transmission — patience, judgement, team culture.

First role (20-22): The engineer arrives operational on the physical layer, with an already-constituted human network, an already-integrated team culture, and a Maâllem who can still answer the phone when something incomprehensible happens at 3am.

◆ WHAT COMPANIES GAIN BY INVESTING IN MENTORSHIP

A less risky hire — recruiting someone already known. Faster integration — the Metaâllem knows the culture, the tools, the people. Stronger loyalty — you do not easily leave a team with which you have a transmission bond. And the perpetuation of tacit knowledge — today's Metaâllem will become tomorrow's Maâllem.

HUMAN

The Zellij of Fès did not survive because there were books about Zellij. It survived because there were Maâllems who accepted taking Metaâllems beside them. Bare-metal infrastructure deserves the same honour.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE GREAT RECHARGE

Hot springs · Mountains · Total freedom

What science says about rest without a programme

HUMAN

◆ **CONTEXT OF THE STUDY**

This study extends "The Body First" on the angle of the great annual recharge — the total disconnection retreat. It argues that a mountain weekend at hot springs, with no imposed programme, no team-building forms, no management debrief, produces a deep cognitive and physiological restoration that no formatted seminar can reach — at a cost accessible to any SME.

Amine RAITI — Infrastructure Architect & SRE

Teaching since 2006 · Electricity to Kubernetes · All audiences

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE PROGRAMME — ITS STRENGTH IS IN WHAT IT DOES NOT HAVE

A WEEKEND WITHOUT AN AGENDA — THAT IS THE INNOVATION

The great annual recharge is not a seminar with a lighter programme. It is a full company retreat — or by session if production or numbers require — with one rule: no imposed programme. The format's strength lies precisely in what it does not have.

◆ WHAT THE WEEKEND CONTAINS

The location: a thermal spa resort in the mountains — accessible by train. Altitude and nature are the two conditions. Everything else is free.

Accommodation: a hotel or gîte taken over by the team. No imposed double rooms, no luxury individual suites — decent rooms, the essentials, the rest of the budget goes to the experience.

Meals: at the company's expense, at whatever time suits each person. No collective lunch at 12:30 sharp. No gala dinner with a CEO speech. Eat when hungry, with whom you want, at the hotel restaurant or the village café.

Activities: thermal cure, bathing in hot springs, hiking, walking, reading, sleeping, massage, nothing at all. Each person does what recharges them. There is no right answer.

What is not there: no team-building workshops, no forms to fill in, no collective feedback sessions, no PowerPoint on company strategy, no external facilitator. None of that. Just free time in the mountains, at the company's expense.

◆ THE BUDGET — ACCESSIBLE TO ANY SME

For a team of 10 people, 2-day weekend:
— Transport (train or van): €50 to €100/person return
— Accommodation (2 nights, double room): €80 to €120/person/night → €160 to €240
— Thermal spa or hot springs access: €20 to €40/person
— Meals (all meals included): €40 to €60/person/day → €80 to €120
— **Total: €350 to €500 per person for 2 days.**
Comparison: formatted corporate two-day seminar = €800 to €2,000 per person. The Great Recharge costs 3 to 4 times less. It produces documented deep neurological restoration. The formatted seminar produces fatigue and post-it notes.

Amine RAITI · CC BY-NC-SA 4.0

WHY MOUNTAINS RESTORE WHAT THE OFFICE HAS EXHAUSTED

Two major scientific theories document why natural environments restore cognitive and emotional resources. They are not contradictory — they describe two complementary mechanisms of the same phenomenon.

◆ ATTENTION RESTORATION THEORY — KAPLAN AND KAPLAN (1989)

Rachel and Stephen Kaplan established that the capacity to focus attention — "directed attention" — is a limited cognitive resource that is exhausted by sustained work. An SRE engineer under permanent tension accumulates attentional fatigue that reduces performance, creativity and the capacity to solve complex problems. Nature restores this resource through "soft fascination" — natural elements capture attention involuntarily and effortlessly (a stream, clouds, snow-covered ridges), freeing directed attention from all demands and allowing it to recover. The four properties of a restorative environment: *Being away* (break from habitual activities), *Extent* (sense of total immersion in a different environment), *Soft fascination* (attention captured effortlessly), *Compatibility* (the place matches what the person seeks). A free mountain weekend at hot springs ticks all four.

◆ STRESS REDUCTION THEORY — ULRICH (1983)

Roger Ulrich established that natural environments trigger evolutionary safety responses — ancestral signals telling the organism it is not in danger. These responses produce positive emotions and reduce stress, measurable by physiological indicators: heart rate, blood pressure, cortisol. Ulrich demonstrated that stressed subjects recover faster when watching nature videos than urban scenes. Cognitive benefits from nature are, in his view, the consequence of emotional and physiological improvement — both theories (ART and SRT) describe distinct but interconnected processes through the same mind-body system, linked by the vagus nerve.

◆ NASSIHA — DOSE MATTERS

The Bell et al. meta-analysis (2025) documents a dose-response relationship between nature exposure and attentional restoration — with diminishing returns beyond a certain threshold. A full weekend in the mountains falls within the optimal return zone — long enough for deep restoration, not so long that the effect plateaus. A single day is insufficient for deep recovery. A month would be too long to be regularly accessible.

HIMAN

Amine RAITI · CC BY-NC-SA 4.0

HOT WATER UNDOES WHAT STRESS HAS DONE

Immersion in thermal hot springs produces documented physiological effects that complement and amplify the restorative effects of nature. Hot springs are not a reward — they are an open-air therapeutic intervention.

◆ WHAT HOT WATER DOES TO CORTISOL

Immersion in water between 37°C and 40°C activates the parasympathetic nervous system — the "rest and digest" mode that opposes the "fight or flight" mode chronically activated in SRE teams under pressure. Cortisol drops. A systematic review published in the *International Journal of Biometeorology* documents that balneotherapy and thermal therapy show potential to influence cortisol levels. Chromogranin — another biological stress marker — is also reduced, more pronouncedly in people with initially higher stress levels. Chronically pressured SRE engineers benefit proportionally more from thermal bathing than low-stress profiles.

◆ VASODILATION AND OXYGENATION

The water's heat causes vasodilation — widening of blood vessels — improving circulation and tissue oxygenation. This improvement in cerebral circulation has documented effects on cognitive clarity and mental fatigue reduction. The minerals present in thermal springs — sulphur, magnesium, bicarbonates — partially penetrate the heat-dilated skin and reinforce these anti-inflammatory and relaxing effects.

◆ SLEEP AFTER THE SPRINGS

A particularly well-documented mechanism: after hot water immersion, body temperature drops naturally upon exit. This thermal drop triggers sleep-inducing processes — the same mechanism as evening temperature drop in temperate environments. Participants in thermal stays report significant sleep quality improvement from the first night. For an SRE team whose members accumulate chronic sleep deficit from on-call duties, this sleep improvement is in itself a workplace health intervention.

HUMAN

◆ THE LINK WITH "THE PAGER AND THE BODY"

This study completes "The Pager and the Body" by proposing its structural antidote. The Pager and the Body documented nocturnal cortisol accumulation from on-call duties. The Great Recharge proposes the intervention that measurably repays this cortisol debt — thermal springs, nature, quality sleep, two days without alerts. This is not a luxury. It is preventive occupational medicine.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · TOTAL FREEDOM — THE KEY MECHANISM

WHY THE ABSENCE OF A PROGRAMME IS THE PROGRAMME

The most counter-intuitive detail of the Great Recharge is also the most important: there is no programme. No planned activities. No mandatory time slots. No forms. This absence is not an organisational oversight — it is the central neurological mechanism of restoration.

◆ PERCEIVED AUTONOMY REDUCES CORTISOL INDEPENDENTLY OF ACTIVITY

Wellbeing psychology research (Deci and Ryan, Self-Determination Theory) establishes that perceived autonomy — the feeling of controlling one's own choices — is one of the most powerful conditions for psychological recovery. A seminar with an imposed programme, however pleasant, maintains a form of directed vigilance: I must be in the right place at the right time, do the right activity, participate correctly. This vigilance keeps the sympathetic nervous system on alert. The total absence of programme frees the brain from this constraint. It can, for the first time in a long while, freely decide the tempo of its own recovery.

◆ FREE HIKING VS GUIDED HIKING

A group that spontaneously decides to hike together does something neurochemically different from a group following a planned guide. In the first case: spontaneous collective decision, pace adjusted to each person, possibility to stop at will, conversations arising naturally from movement. In the second: obligation to follow, imposed pace, attention partially directed toward the guide. Free hiking produces Kaplan's ART effect — soft fascination, being away, extent. Guided hiking produces a collective activity with slight performance constraints.

◆ EAT WHEN YOU WANT — THE TRUST SIGNAL

Letting team members eat when it suits them — alone, in spontaneous small groups, at the hotel restaurant or the village café — is a management trust signal that reinforces psychological safety. Nobody is monitored. Nobody is assessed. This explicit trust signal is one of the mechanisms documented by Amy Edmondson as a condition for collective high performance — and it is produced here for free, by a simple organisational decision.

◆ NASSIHA — CHOSEN SOLITUDE IS A RECOVERY ACTIVITY

Some team members will choose to spend part of the weekend alone — reading in their room, walking solo, staying longer in the springs. This chosen solitude is not a cohesion failure — it is a legitimate form of recovery that the format must explicitly protect. A seminar that leaves no room for chosen solitude is a seminar that does not understand introverts.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · VACATION RESEARCH — WHAT DISCONNECTION PRODUCES

THE MEASURABLE EFFECTS OF A REAL REST WEEKEND

Research on vacation effects on workplace performance is convergent and robust. It documents real, measurable benefits, but also their limits — which allows optimising the Great Recharge format to maximise lasting effects.

◆ WHAT DE BLOOM ET AL. META-ANALYSIS ESTABLISHES

The de Bloom et al. meta-analysis (22 studies) documents significant reductions in exhaustion, improvements in mood and life satisfaction after vacations. Grant et al. (54 studies) identify measurable physiological improvements: cortisol reduction, better heart rate variability — an indicator of autonomic nervous system recovery — and better sleep quality. These effects are documented from two days of real disconnection.

◆ REAL DISCONNECTION — A NECESSARY CONDITION

The documented effects require real disconnection — not a weekend where team members check their Slack alerts every two hours. The management decision to guarantee minimal on-call during the weekend — with a degraded incident management protocol, without pressure on the present team — is the organisational condition for the neurological benefit. If engineers go to the mountains but remain responsible for production, cortisol does not drop. Disconnection is binary: one is in recovery or one is not.

◆ THE FADE-OUT EFFECT AND ITS RESPONSE

Vacation research also documents a "fade-out" — benefits progressively diminish in the weeks after return. The response to this fade-out is not to extend the annual retreat — it is to complement it with more frequent, less intensive formats. This is exactly the architecture proposed in this corpus: Great Recharge once a year, quarterly team dinner, seasonal hammam + walk, weekly in-office play. Recovery is continuous maintenance, not an annual deep clean.

◆ EFFORT-RECOVERY THEORY

Effort-Recovery Theory (Meijman and Mulder, 1998) establishes that sustained work effort depletes physiological and psychological resources — elevated cortisol, cognitive fatigue, emotional exhaustion. Without periodic recovery, this load accumulates and increases burnout risk. Vacations interrupt this cycle. The mountain Great Recharge is precisely the recovery structure this theory calls for.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL — THE COMPLETE RECHARGE ARCHITECTURE

FROM THE WEEKLY EVENING TO THE ANNUAL WEEKEND

This study completes the recharge architecture proposed in the corpus. Each format has a frequency, a cost and a distinct neurochemical mechanism. Together, they cover the full spectrum of recovery — from micro-daily to the great annual.

◆ **THE COMPLETE ARCHITECTURE — FREQUENCY, COST, MECHANISM**

Daily / Weekly — Play in SRE (Nerf, active break): €0 · Micro adrenaline discharge, fiero, dehierarchisation · 10 minutes

Monthly — Simple team dinner: €30 to €78/person · Bonding oxytocin, shared memory, dehierarchisation · 4 hours

Quarterly — Hammam + free walk + dinner: €35/person · Cortisol ↓, oxytocin ↑, endorphins, natural sub-groups · 4 hours

Biannual — Competitive activity (karting, bowling, opera): €50 to €100/person · Fiero, adrenaline, collective emotion · 3 to 4 hours

Annual — The Great Recharge (mountains, hot springs): €350 to €500/person · Deep attentional restoration (ART/SRT), cortisol ↓↓, sleep, autonomy, informal bonding · 2 days

Total annual budget per person: €600 to €900.

The cost of a single formatted seminar day — for a complete recovery architecture covering all neurochemical levels, from the weekly micro-dose to the deep annual restoration.

◆ **THE MANAGEMENT DECISION THAT PRECEDES EVERYTHING**

The Great Recharge is not decided the day before departure. It is planned 2 to 3 months in advance — to guarantee production coverage during the absence, to allow everyone to organise their personal life, and so that anticipation itself contributes to recovery. Research documents that anticipation of a positive event produces beneficial effects on mood from the announcement — the well-being peak sometimes precedes the weekend itself.

◆◆◆

A mountain weekend at hot springs. No programme. No forms. No consultant. Science says it restores what on-call has exhausted. The budget says it is accessible. What is missing is the decision.

◆◆◆

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE STATE AND THE INVISIBLE BODY

Why digital sovereignty begins by protecting its engineers

RATIO

◆ CONTEXT OF THE STUDY

The French state regulates and protects the professions whose disappearance would threaten the continuity of social life. It does not regulate — and does not protect — the body of bare-metal digital infrastructure engineers and technicians. This study poses the question of this asymmetry and formulates six concrete public policy levers to correct a situation that directly threatens French and European digital sovereignty.

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE QUESTION THAT ADMITS NO COMFORTABLE ANSWER

DOES THE STATE TOLERATE THE DISAPPEARANCE OF OTHER STRATEGIC BODIES?

The question is simple. It deserves a direct answer before any analytical development.

Does the state tolerate the disappearance of the medical profession? No. It created the Ordre des médecins, the numerus clausus, the status of hospital practitioner, conventions with the health insurance system. When medical deserts appeared, it declared a national emergency and mobilised billions.

Does the state tolerate the disappearance of pilots? No. EASA licences are regulated, qualifications maintained by legal obligation, initial training funded and supervised.

Does the state tolerate the disappearance of the construction body? No. Qualibat, Qualifelec, RGE, ten-year liability — the construction body is certified, regulated, protected by legal obligations binding on project owners.

Does the state tolerate the disappearance of notaries, pharmacists, architects? No. Regulated professions, controlled access, protected titles, criminal sanctions for unlicensed practice.

◆ THE QUESTION THAT FOLLOWS

Why then does the state tolerate the slow assassination of the only professional body capable of leading the digital sovereignty project?

The infrastructure technician, the systems administrator, the bare-metal SRE engineer — the person who can cable a switch, diagnose a network fault at 3am, build a sovereign infrastructure outside the hyperscaler cloud — has no protection. No protected title. No enforceable skills framework. No continuing education obligation. Anyone can declare themselves a "DevOps engineer" tomorrow morning and nobody can legally object.

Without this body, OVHcloud, Ecritel, Scaleway and their European equivalents will no longer find staff for their server rooms. Digital sovereignty is not a question of contracts — it is a question of hands and competences. And those hands are disappearing.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

THE DOCUMENTED DISSOLUTION

The Opération Dindon corpus documented, study after study, the mechanisms of the dissolution of the infrastructure engineering body. This is not an intuition — it is a structural observation accumulated across twelve studies and several years of field observation.

◆ MECHANISM 1 — COGNITIVE CAPTURE BY HYPERSCALERS

"Anatomy of the Loss" documented how cloud certifications replaced infrastructure training in skills frameworks. A certified AWS engineer who has never seen a server room is hired instead of a technician who has maintained systems for ten years. The market follows certifications. Certifications follow hyperscalers. Hyperscalers have every interest in the market depending on them.

◆ MECHANISM 2 — DILUTION BY TITLES

"The Infrational Crisis" documented how "Infrastructure as Code" and "DevOps" emptied words of their substance. The bare-metal SRE becomes "DevOps engineer". The systems administrator becomes "cloud engineer". Titles change, physical layer competences disappear from job descriptions, and HR recruits on the new titles without understanding what was lost in conversion.

◆ MECHANISM 3 — THE CONFERENCES THAT LAMENT WITHOUT ACTING

There exist in France and Europe dozens of associations, think tanks, conferences and working groups on "digital sovereignty". They publish reports. They organise roundtables. They produce recommendations. Structurally, they change nothing. Digital sovereignty is not decreed in a conference room — it is built in server rooms, by trained engineers kept in activity. Without this body, all the reports in the world will not produce a functional sovereign datacentre.

◆ THE PREDICTABLE RESULT IN 10 YEARS

If the current trend continues without state intervention, OVHcloud, Ecritel, Scaleway and their equivalents will no longer find the competences needed to operate their infrastructures. They will be forced to either massively import skills from third countries, or sub-contract their operation to providers, some of whom will be foreign. Digital sovereignty will have been lost not by political decision but by the silent exhaustion of a professional body that nobody thought to protect.

Amine RAITI · CC BY-NC-SA 4.0

THE MODELS THAT WORK — IN OTHER BODIES

The state is not powerless. It knows how to protect professional bodies when it decides their disappearance is unacceptable. The models exist — they simply need transposing.

◆ MODEL 1 — THE REGULATED PROFESSION WITH PROTECTED TITLE

Doctor, pilot, architect, notary — these titles are protected by law. Practising them without holding them is a criminal offence. This protection creates market pressure toward quality: clients wanting these competences must use qualified professionals. The same mechanism applied to "critical infrastructure engineer" or "certified systems and networks technician" would produce an immediate revaluation of the body.

◆ MODEL 2 — MANDATORY CERTIFICATION IN PUBLIC PROCUREMENT

To build a public building, a contractor must present certifications (Qualibat, Qualifelec). To participate in a public aeronautical maintenance contract, DGAC approvals are required. Why is a provider operating the infrastructure of a hospital, a town hall, or a ministry not required to present equivalent infrastructure competence certifications? The absence of this obligation in digital public procurement is a gaping hole in sovereignty policy.

◆ MODEL 3 — PROTECTED AND FUNDED INITIAL TRAINING

BTS Digital Systems, BTS SIO, DUT Networks & Telecoms — these training programmes exist. They are underfunded, undervalued in school guidance, and progressively cloud-oriented at the expense of the physical layer. The state that decides infrastructure training is strategic injects resources, revalues the pathways, and explicitly integrates bare-metal into the frameworks — as it integrated aeronautical mechanics into aeronautical technical training.

◆ NASSIHA — REGULATION MUST NOT CLOSE DOORS TO SELF-TAUGHT PROFILES

The corpus documented in "The Blue" the risk of regulation that excludes self-taught profiles. Any infrastructure body protection policy must integrate access pathways through practice — VAE, permanent IDPE, experience recognition — and not reserve access solely to holders of formal diplomas. The protection of the body must protect competences, not pathways.

Amine RAITI · CC BY-NC-SA 4.0

THE COST OF INACTION — STRUCTURAL CALCULATION

The inaction argument is always the same: "protecting this body costs money, creates rigidities, hinders innovation." This argument ignores the cost of inaction — which is systematically higher than the cost of preventive action, in every domain where the state has learned this the hard way.

◆ COST 1 — INDEFINITE HYPERSCALER DEPENDENCE

Every infrastructure engineer who leaves the market to retrain as a "cloud engineer" is an engineer who will not return to bare-metal. France and Europe's dependence on American hyperscalers deepens in direct proportion to the shortage of sovereign infrastructure competences. This dependence has an economic cost (egress fees, contractual lock-in, hyperscaler pricing power) and a geopolitical cost (CLOUD Act, legal extraterritoriality, vulnerability to unilateral American decisions) that the state bears without ever calculating it explicitly.

◆ COST 2 — THE OPERATIONAL INCAPACITY OF SOVEREIGN ACTORS

OVHcloud, Ecritel, Scaleway, Infomaniak, Clever Cloud — these actors exist. They have datacentres. They have offerings. They do not have enough qualified staff to operate them at the necessary scale. This is not a problem of commercial will — it is a problem of competence pool. If the state decided tomorrow that all public administrations must host their data with a European sovereign actor, these actors could not absorb the demand for want of trained infrastructure technicians and engineers. A preference policy only makes sense if the body to operationalise it exists.

◆ COST 3 — NATIONAL SECURITY

National critical infrastructures — energy networks, telecommunications networks, hospital systems, transport networks — run on computing systems that infrastructure engineers maintain. The shortage of these engineers is a national security vulnerability that neither CISOs nor cybersecurity consultants can compensate. A CISO without an infrastructure engineer to implement their recommendations is a doctor without a nurse to administer the care.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHY THE MARKET DOES NOT CORRECT — AND WILL NOT

The classic liberal argument in the face of a skills shortage is that the market self-corrects: if infrastructure engineers are scarce, their salaries rise, attracting new entrants, resolving the shortage. This argument is sound in an equilibrium market. It does not apply when powerful actors have an interest in equilibrium not being reached.

◆ WHY HYPERSCALERS HAVE NO INTEREST IN THE BODY RECONSTITUTING

A market rich in qualified infrastructure engineers is a market where clients can choose not to use hyperscalers. The shortage of bare-metal competences is structurally favourable to hyperscalers — it makes cloud even more unavoidable because there are no longer the competences to do otherwise. Hyperscalers do not actively train bare-metal engineers. They train "cloud architects" certified on their own platforms. This is not malevolence — it is economic rationality.

◆ WHY INDIVIDUAL EMPLOYERS DO NOT INVEST

Training an infrastructure engineer takes 2 to 5 years. The employer who invests in this training runs the risk that the engineer leaves for a competitor or a hyperscaler once trained. This free-riding risk reduces private investment in infrastructure training below the socially optimal level — exactly the same mechanism that justifies state intervention in the training of doctors, pilots and nuclear engineers.

◆ WHY CONFERENCES ARE NOT ENOUGH

Digital sovereignty professional associations and conferences produce visibility and networking. They do not produce technicians. Visibility without a trained professional body is a conference on medical desertification without training doctors. Speech is necessary. It is not sufficient. And when it becomes the substitute for structural action, it is counterproductive — it gives the impression that something is being done when nothing is.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL — SIX LEVERS

WHAT THE STATE CAN DO — NOW

◆ LEVER 1 — CREATE A FORMAL "CRITICAL INFRASTRUCTURE ENGINEER" STATUS

A title registered in the RNCP, accessible by curriculum and by VAE, recognised in collective agreements with a protected salary scale. This title protects competence — not the pathway. It is accessible to the self-taught Blue through VAE as much as to the engineering school graduate. It creates market pressure toward recognition of bare-metal competences.

◆ LEVER 2 — MANDATORY CERTIFICATIONS IN DIGITAL PUBLIC PROCUREMENT

Any provider operating the infrastructure of a public body — administration, hospital, local authority — must demonstrate engineers holding a recognised infrastructure competence title. Like Qualibat for construction. Like DGAC approvals for aeronautical maintenance. This obligation creates a market for competence and a competitive disadvantage for providers who have substituted cloud generalists for infrastructure engineers.

◆ LEVER 3 — REINTEGRATE BARE-METAL INTO BTS AND DUT FRAMEWORKS

BTS SIO, BTS Digital Systems, DUT Networks & Telecoms must explicitly integrate the physical layer into their frameworks: cabling, hardware, bare-metal systems, physical network diagnosis. The Foundation of Iron — 26 weeks from electricity to networking — is a turnkey programme available under an open licence. It can be adapted in 12 months into a 2-year BTS curriculum.

◆ LEVER 4 — CONVENTION THE FOUNDATION OF IRON WITH FRANCE TRAVAIL

The historical pathway exists and worked for fifteen years. France Travail directly funds accredited training organisations on shortage-area training programmes. Bare-metal infrastructure is a documented structural shortage area. The AFC convention is sufficient — no legislation needed, no exceptional budget. An administrative decision.

◆ LEVER 5 — EUROPEAN PREFERENCE IN DIGITAL PUBLIC PROCUREMENT

Integrate an explicit preference for sovereign European actors (OVHcloud, Ecritel, Scaleway, Hetzner, Infomaniak) into the award criteria of digital public procurement. This preference creates infrastructure employment in France, funds sovereign actors, and reduces dependence on the trio. A state that preaches digital sovereignty while hosting its data on AWS is not practising sovereignty — it is talking about it.

◆ LEVER 6 — ANNUAL REPORT ON THE STATE OF THE INFRASTRUCTURE BODY

The state publishes each year the medical demographics, the state of medical deserts, training needs. It publishes nothing equivalent on digital infrastructure competences. Creating this reporting obligation — number of active SRE engineers and systems administrators, evolution over 5 years, cloud reconversion rate, shortages by region — makes visible what is invisible. And what is not measured is not protected.

◆◆◆

The state does not tolerate the disappearance of doctors. It must not tolerate the disappearance of the engineers who maintain the nation's digital infrastructure. These are strategic bodies. They deserve the same treatment.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE UNIFORM OF THE BODY

Making the Invisible Visible

The Dress Code of Infrastructure

HUMAN

◆ POSITION OF THIS STUDY

The uniform is the physical manifestation of the SysOps/NetOps/OpInfra nomenclature documented in "The Ticket and the Talent". It makes visible what has until now been purely administrative. It materialises the body — exactly what the corpus has defended since 8 May 2026. A body without a uniform is a body without boundaries: difficult to defend, difficult to value, difficult to transmit.

LEVELS	COSTS	WATERMARK
5	**EMPLOYER**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · WHY A UNIFORM IN INFRASTRUCTURE

MAKING THE INVISIBLE VISIBLE

Infrastructure is the only technical body without a visible distinguishing mark. The doctor wears the white coat. The electrician wears the helmet and vest. The chemist wears the lab suit. The SRE engineer who keeps alive the systems on which the entire organisation depends arrives at the office in jeans and leaves in jeans. Nobody knows who they are, what they do, or what level of responsibility they carry.

◆ THE INVISIBLE BODY — WHAT THE ABSENCE OF UNIFORM SAYS

The absence of uniform says three things simultaneously. It says that this body's work is no different from others' work — which is false. It says that levels of responsibility are interchangeable — which is dangerous. And it says that contact with the machine requires no particular protection or signalling — which is a profound cultural error.

In every technical profession that has built its professional dignity, the uniform preceded or accompanied institutional recognition. The doctor's white coat dates from the 19th century — it preceded professional orders, licences, statutes. It said: this person is in a professional relationship with a subject that deserves the sign of that relationship. The datacentre deserves the same sign.

◆ THE LINK WITH THE TICKET AND THE TALENT

"The Ticket and the Talent" proposed a structured nomenclature of infrastructure roles: three categories (SysOps/NetOps/OpInfra), six levels, the Non-Dilution Rule. This nomenclature exists on paper. It does not exist in physical space — in the meeting room, in the datacentre, in the workshop. The uniform is the physical materialisation of this nomenclature. It says, without a word, what level of responsibility one is dealing with. It makes the technical hierarchy legible for those who do not know it — and respectable for those who operate within it.

◆ NASSIHA — THE UNIFORM IS NOT A CONSTRAINT

The first objection will be: "We are not a factory." The right answer is: the doctor is not in a factory either. The uniform is not a mark of subordination — it is a mark of belonging to a recognised professional body. The difference is fundamental.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

FIVE LEVELS · FIVE SIGNALS · ONE COHERENCE

LEVEL	UNIFORM	CONTEXT
Infrastructure Architect	Civilian dress — personal choice	Meetings, design, steering. Rarely on the physical floor.
Lead SysOps · Lead NetOps · Principal SRE	♂ **Black suit · White shirt · Red tie** ♀ Three equivalent options — personal choice: ① Black trouser suit · White shirt · Red scarf ② Black blazer · White collar · Red brooch ③ Black sheath dress · White collar · Fine red scarf	Operational and strategic. Strong signal of responsibility level. All three feminine options are equivalent in authority — same black/white/red triptych, freedom of form.
Engineer · Administrator	**White coat** (field) · Civilian (office/remote)	Coat mandatory in datacentre, workshop, lab. Civilian in administrative context.
Field Technician	**Khaki green coat** (military field)	Permanent field work. Robustness, functionality, technical connotation.
Support & Proximity Technician	**Blue coat**	User and desktop equipment contact. Signal of availability and service.

◆ JUSTIFICATION OF THE BLACK/WHITE/RED TRIPTYCH FOR LEADS

The Lead SRE or Principal SRE is the highest level of operational responsibility on the floor. He — or she — makes decisions during major incidents, arbitrates priorities, engages the body's responsibility on technical choices. This level deserves a strong visual signal — not jeans. Red is the signal of operational command: in all structured bodies, this colour signals decision-making authority. Black says: this is not a technician, not a manager — this is a technical commander.

Women Leads have three equivalent options — all within the black/white/red triptych, all at the same level of visual authority, with a freedom of form that the male suit does not offer. This is not a concession — it is a recognition that feminine elegance has more registers, and that the infrastructure body must welcome them all. A dress code that had not considered its feminine expression would reproduce exactly the Invisible Amputation this corpus has documented.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE WHITE COAT AS PROFESSIONAL ACT

THE DATACENTRE AS A PLACE OF CARE — THE ENGINEER AS PRACTITIONER

The white coat is not protective clothing in the PPE sense — it is a professional signal. It says: I am in a professional relationship with what I am touching. This relationship demands method, precision, responsibility. The engineer who enters the datacentre in a white coat says something that the one who enters in jeans does not.

◆ THE MEDICAL METAPHOR — APPLICABLE WITHOUT RESTRICTION

The doctor wears the white coat because they enter into a care relationship with a patient. The coat says: what I do here is serious, methodical, and engages my professional responsibility. The engineer who opens a production rack enters the same type of relationship with the system on which users, data, and sometimes lives depend. There is no reason this relationship should be less signalled than the doctor's with their patient.

The infrastructure engineer's white coat says: I am in a professional relationship with a critical infrastructure. What I do here can affect thousands of people. I know this. My uniform says I know this.

◆ THE OBLIGATION CONTEXT — FIELD ONLY

The white coat is mandatory only in physical field contexts: datacentre, workshop, lab, server room, physical test environment. It is not required at the office, while working remotely, or in meetings. This distinction is fundamental — it says the uniform is not a bureaucratic constraint, it is an act tied to contact with the machine. When the engineer leaves the server room, they remove the coat. Like the doctor leaving the operating theatre.

◆ THE COAT AS SYMBOLIC PROTECTION

In advanced technical disciplines, the uniform also protects its wearer in another way: it protects against ambiguity of role. The engineer in a white coat in the datacentre is not approached for user support. They are not confused with an external contractor. Their signal says clearly: I am the practitioner of this machine. The coat creates a visible professional boundary that protects concentration, method and responsibility.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · COSTS BORNE BY THE EMPLOYER

THE ECONOMIC CALCULATION — AND WHAT IT REVEALS

◆ COST ESTIMATE BY LEVEL

Lead/Principal — male (full black suit): 2 suits × €300-500 = €600-1,000 acquisition. Dry cleaning: €15/suit × 24 times/year = €360/year. Annual total: €560-700/person (amortised over 3 years).

Lead/Principal — female option ① (black trouser suit): 2 suits × €250-450 = €500-900 acquisition. Dry cleaning: €12 × 24 times/year = €290/year. Annual total: €460-600/person.

Lead/Principal — female option ② (black blazer + brooch): 2 outfits × €200-350 + brooch €20-80 = €420-780 acquisition. Maintenance: €240/year. Annual total: €380-500/person.

Lead/Principal — female option ③ (sheath dress + scarf): 2 dresses × €150-300 + 2 scarves × €30-80 = €360-760 acquisition. Laundry: €200/year. Annual total: €320-450/person.

Combined range Lead/Principal — all options: €320-700/year/person — equivalent across male and female options, below the cost of one AWS cloud certification (€2,300-4,300).

Engineer/Administrator (white coats): 3 coats × €30-50 = €90-150 acquisition. Laundry: €3/coat × 48 weeks = €430/year. Annual total: €460-510/person.

Field Technician (khaki green coats): 3 coats × €25-40 = €75-120 acquisition. Laundry: €3 × 48 weeks = €430/year. Annual total: €455-490/person.

Support Technician (blue coats): Same structure. Annual total: €455-490/person.

Architect: €0 (personal civilian dress, no cost to employer).

◆ COMPARISON WITH CURRENT EXPENDITURE

An AWS Solutions Architect Professional certification costs €300 in exam fees, plus €2,000-4,000 in training. Total: €2,300-4,300 per person. A Google Cloud Professional: €200 exam + €1,500-3,000 training. Organisations spend without hesitation €2,000 to €4,000 per engineer on cloud certifications that reinforce hyperscaler lock-in. They hesitate before €500/year for a uniform that reinforces the identity and dignity of the body. This contrast says something about organisational priorities.

◆ THE UNIFORM AS AN HR BENEFIT

Laundry and maintenance costs borne by the employer constitute a net salary benefit. For an engineer whose coat is laundered by the employer, this represents €430/year of preserved purchasing power — non-taxable if treated as professional expenses. In a context of documented salary compression described in "The Infrational Loop", this is an immediately activable lever without revising the pay scale.

Amine RAITI · CC BY-NC-SA 4.0

THREE OBJECTIONS · THREE RESPONSES

◆ OBJECTION 1 — "WE ARE NOT A FACTORY"

The objection: Infrastructure is not an assembly line. Our engineers are knowledge workers, not manual operators. The uniform degrades their status.

The response: The doctor is not in a factory. The pharmacist is not in a factory. The laboratory researcher is not in a factory. All wear a uniform because they practise a technical profession in a context that requires a signal of competence and responsibility. "Knowledge worker" and "wears a uniform" are not contradictory — they are complementary. The uniform does not degrade status — it signals it.

◆ OBJECTION 2 — "IT WILL DRIVE TALENT AWAY"

The objection: SRE engineers are rare and demanding about their work environment. Imposing a uniform will push them toward freer organisations.

The response: The proposed uniform is not permanently imposed — it is contextual. The white coat only in physical field contexts, the suit for Leads only in their technical command role. Outside these contexts, civilian dress remains the norm. Moreover, the engineer who chooses an organisation because it visually recognises their professional body is an engineer who identifies with that body. This is a retention signal, not a repulsion signal.

◆ OBJECTION 3 — "IT IS NOT OUR CULTURE"

The objection: French and European tech culture values informality. Suits and uniforms belong to banking and the military, not digital infrastructure.

The response: Tech culture also explains why infrastructure is invisible, undervalued and underpaid. If the current culture produces the Infrational Loop, the Invisible Amputation and the Departure of the Last One Who Knows — then this culture must evolve. The uniform is not an import from another culture: it is the construction of a culture specific to the infrastructure body, coherent with its actual responsibility.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE UNIFORM AS AN ACT OF SOVEREIGNTY

A BODY WITHOUT BOUNDARIES IS A BODY WITHOUT DEFENCE

The corpus's founding thesis — "No Sovereignty Without Matter" — applies to the professional body too. No sovereignty of the body without visible boundaries. A body whose members are indistinguishable from other workers in the organisation is a body without its own territory — impossible to defend, impossible to value, impossible to transmit.

◆ THE UNIFORM AS THE VISIBLE SKIN OF TECHNICAL SOVEREIGNTY

Technical sovereignty rests on three conditions documented in this study and the corpus: engineers who know, tools that enable, political decisions that create conditions. The uniform is the fourth condition — the one that makes the first three visible. An engineer without a uniform in a datacentre is knowledge without a signal. A Lead SRE without a suit in a crisis meeting is authority without a mark. Technical sovereignty begins in bodies — and bodies must be legible.

◆ THE COMPLETE MODEL — NOMENCLATURE + UNIFORM + TRAINING

The corpus has built a coherent three-layer model. "The Ticket and the Talent" gave the nomenclature — SysOps/NetOps/OpInfra, six levels, Non-Dilution Rule. "The Foundation of Iron" gave the training — 26 weeks, electricity to Kubernetes. "The Uniform of the Body" gives visibility — five levels, five signals, one coherent code. These three layers together constitute a complete professional body: named, trained, visible. This is what no French organisation has yet built for its infrastructure.

◆ WHAT THE UNIFORM SAYS THAT THE CONTRACT DOES NOT

The employment contract states the hierarchical level. The payslip states the salary. The badge states the company. The uniform says what the other three do not: which professional body one belongs to, what relationship one has with the machine, what level of technical responsibility one carries. It is a signal that precedes speech, that crosses the meeting room before the engineer has opened their mouth. In a world where the infrastructure body is invisible, this signal is revolutionary.

HUMAN

◆◆◆

A body without a uniform is a body without boundaries.
A body without boundaries is a body without defence.
A body without defence is a body that disappears.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE COAT AS SOCIAL LIFT

THE SIGN THAT PRECEDES PROMOTION — THE STRIPE THAT MOTIVATES

This is the angle the previous six sections did not address — and the most powerful one. The coat is not only a signal of belonging. It is a visible, daily, physical social lift. The one in blue will want khaki. The one in khaki will want white. The one in white will want the black suit — or the trouser suit, or the sheath dress. This desire is not abstract — it is written into the fabric of the break room, the corridor, the datacentre.

◆ THE MILITARY MODEL — THE RANK AS PHYSICAL MOTIVATION

The military rank is the most documented model of social mobility through visible sign. The soldier looking at the sergeant's stripes does not see an HR abstraction — they see a physical objective, dated, attainable through known criteria. It is not "develop your career" — it is: these stripes exist, these criteria are public, this progression is possible. The motivation is daily because the sign is daily.

The support technician in blue coat crossing the engineer in white coat in the datacentre corridor sees exactly the same thing. Not a promise in an annual review — a physical object, worn by a real person, saying: this level exists, it is attainable. The white coat is more motivating than a pay rise promised eighteen months from now.

◆ TRANSPARENCY OF CRITERIA — THE COAT OBJECTIVELY EARNED

The lift only works if the criteria for changing coat are public, objective and uncontestable. The coat does not change because the manager decided — it changes because the nomenclature criteria documented in "The Ticket and the Talent" are met. Six levels, three categories, defined competence criteria. Nobody can refuse the white coat to the one who has the criteria. Nobody can take it from the one who maintains them.

This transparency protects against favouritism, arbitrariness and discrimination. It says: the path is there, it is the same for everyone, and it is measurable. This is what the annual review with a well-intentioned but subjective manager cannot guarantee.

◆ THE RESPONSIBILITY OF THE SIGN — THE CHAIN OF GAZE

The mechanism works in both directions. The one wearing red — tie, scarf, brooch or sash — is watched by those wearing the white coat. The one in white is watched by those in khaki. The one in khaki is watched by those in blue. This chain of gaze is a professional conduct mechanism — not imposed by internal regulations, but by the social weight of the sign. One does not wear the red signal without feeling watched by those who would want to wear it. This gaze is a daily responsibility — more effective than any management by objectives.

◆◆◆

The one in blue will want khaki. The one in khaki will want white.
The one in white will want the black suit — or the trouser suit, or the sheath dress.
This desire is the lift. The sign is the engine.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

PART VII

EDUCATION & TRAINING

7 pieces

HUMAN

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE BLUE

What the diploma does not measure

Self-taught, apprentices, interns

and the case for a permanent IDPE

RATIO

◆ **CONTEXT OF THE STUDY**

This study is part of the Opération Dindon corpus, following "Anatomy of the Loss" (dissolution of infrastructure competence), "The Mastery of Iron" (hardware concentration), "The Economy of Silence" (uptime as the sole SRE metric) and "The Silence of the Rooms" (women's absence from infrastructure). It addresses a fourth blind spot: competent profiles excluded from the market by the diploma and certification filter — and puts forward a concrete public policy proposal to correct this filter, using a mechanism that already exists in French law.

Amine Raiti · Infrastructure Architect & SRE · Former instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE CASE IN POINT

THE BLUE

At a final-year internship interview, a young man walks in. No engineering degree, no vendor certification, no impressive line on a still-empty CV. But within ten minutes, he explains how he assembled and optimised a cryptocurrency mining platform — hardware selected component by component, cooling calculated, power consumption compressed to a minimum because his grandmother is the one paying the electricity bill.

This young man understands electrical power, thermal dissipation, optimisation under real budget constraints, and hardware cost-to-performance ratios. He taught himself everything, on his own money, under a financial discipline that makes many degreed engineers uncomfortable. He is hired that day. He will become an excellent SRE. He is called "the Blue" — because he arrives without colours, and acquires them quickly.

WHAT THE BLUE KNOWS THAT THE CV DOES NOT SAY

Electrical optimisation under real financial constraints is one of the most directly transferable skills to datacentre administration. Someone who reduced a mining platform's power draw by 15% to stay within a family budget has practised, without naming it, physical FinOps. This is not a hobby. It is field engineering.

The Blue is not an isolated case. He represents an entire category of profiles that the infrastructure market systematically excludes before the interview — because their competence does not present itself in the right format. This study analyses the mechanisms of that exclusion and proposes a correction lever that already exists in French law.

◆ NASSIHA — METHODOLOGICAL NOTE

The word "diploma" is used in this study in a broad sense: it denotes the full set of formal titles — state diplomas, RNCP titles, vendor or publisher certifications — used as entry filters in the infrastructure job market. The critique targets the use of these titles as substitutes for assessing real competence, not their existence or pedagogical usefulness.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE PAPER CULTURE

THE DIPLOMA AS A PROXY — WHAT IT FILTERS CORRECTLY AND WHAT IT FILTERS WRONGLY

A diploma is a proxy. It does not directly measure competence — it measures the ability to have completed a defined curriculum, in a defined timeframe, with the resources needed to do so. This proxy is useful in many cases: it guarantees a common foundation, it reduces recruiter uncertainty when facing an unknown profile, it signals structured learning capacity. These utilities are real.

But a proxy is not a direct measure. When the proxy becomes the sole access criterion, it mechanically excludes all profiles whose real competence exceeds what the proxy can capture. In systems and network infrastructure, this decoupling is particularly frequent — because infrastructure competence is acquired massively through practice, through exposure to failures, through autonomous technical curiosity. Vectors that do not run through formal curricula.

◆ WHAT THE DIPLOMA FILTERS CORRECTLY

A verifiable minimum theoretical foundation. Exposure to structured working methods. The ability to learn within an institutional framework. These are useful things to verify, and the diploma does it reasonably well for profiles who followed that path.

◆ WHAT THE DIPLOMA FILTERS WRONGLY

Competence acquired outside formal education — through self-teaching, personal projects, years of production without a safety net. The capacity for diagnosis under pressure, which is not learned in a lecture theatre. Ingenuity under real constraints, which the Blue developed because his grandmother had no unlimited budget. These competences leave no trace in a diploma system — they leave a trace in the systems that keep running.

◆ THE RATIONAL RECRUITER PARADOX

A recruiter who uses the diploma as a hard filter acts rationally at the individual level — they reduce their sorting cost. But at sector level, they collectively produce a market that under-uses its pool of real competence and overpays for titles that do not always correspond to operational capability. The aggregated result of individually rational decisions is collectively suboptimal. This is a classic system-level effect.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

THE INTERN AS A VECTOR OF FRESH KNOWLEDGE

The previous section described the intern as a profile exposed to extraction without return — they contribute, they leave without guarantees. This description is accurate but incomplete. It only captures one direction of the flow. There is a second, more subtle one, systematically ignored in the discourse on internships: the intern also brings something the team in place does not have.

The senior SRE who has been maintaining a production infrastructure for five years is excellent at what they do. They know their system better than anyone. But their R&D time is zero or close to zero. Production absorbs everything. Deep technology updates, emerging architectures, new approaches — they hear about them, they do not practise them. They are up to date on what is running. They are behind on what is coming.

◆ THE BLUE TAUGHT ME SOMETHING

When the Blue joined the team, he introduced concepts nobody had practised: mining hardware, blockchain architecture, power consumption optimisation under constraint. These were not yet in our infrastructures. They were in his garage. He was six months ahead of the technical subjects that the sector would raise as serious questions two years later. His ignorance of how our system worked was real — and temporary. His knowledge of emerging technical ground was a resource we did not have and could not have obtained any other way.

◆ WHY DOES THE INTERN HAVE THIS LEAD?

The intern does not yet carry the constraints of production. They do not have a system to keep alive at 3am. They do not have incident tickets to clear before they can think about anything else. This absence of constraint, which looks like a deficit, is precisely what allows them to spend time on research, experimentation, and exploring technologies that do not yet have an established use case. They explore while the senior maintains. Both are necessary. Neither replaces the other.

◆ A BIDIRECTIONAL FLOW THAT INTERNSHIPS SHOULD FORMALISE

The implicit model of an internship is unidirectional: the company transmits its knowledge to the intern, the intern learns. This model under-exploits the relationship. A well-designed internship explicitly formalises the flow in both directions — the intern learns the production system, the team learns from the intern about what they practise outside production. This is not goodwill. It is knowledge management.

◆ NASSIHA — THE LIMIT OF THIS THESIS

This thesis does not apply uniformly to all interns. It applies to actively self-taught profiles — those with an intense personal practice outside the curriculum, like the Blue. An intern without autonomous practice does not necessarily bring fresh knowledge. The condition for activating this reverse flow is the existence of autonomous technical curiosity in the intern — precisely the profile that current recruitment filters tend not to recognise.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THREE EXCLUDED PROFILES

SELF-TAUGHT · APPRENTICE · INTERN

Three distinct profiles encounter the diploma and certification filter in infrastructure, through different mechanisms. They share one thing in common: their real competence is systematically undervalued by the formal entry filters of the job market.

◆ PROFILE 1 — THE SELF-TAUGHT

The Blue is the archetype. Real technical competence, sometimes surpassing the average graduate, acquired through practice, curiosity and exposure to real problems. His obstacle: he does not pass the automatic filters of ATS (Applicant Tracking Systems) that screen out CVs without a diploma before a human eye reads them. He does not exist for the formal market — even if he has been running production systems for years. His only current lever: finding a recruiter or instructor willing to look at what he can do rather than what he has on paper.

◆ PROFILE 2 — THE APPRENTICE

The apprentice faces the opposite situation: institutional legitimacy in the process of being acquired, but treated by the market throughout their training period as a "not yet ready" profile. They are on the ground, genuinely contributing to production, but their "still a student" status creates friction at the point of conversion to permanent employment — some employers prefer to keep the apprentice as an apprentice rather than converting them, for cost and flexibility reasons. The entry pathway recommended by the sector paradoxically produces its own extraction loops.

◆ PROFILE 3 — THE INTERN

The intern is the profile most exposed to extraction without return. They make a real contribution — often on tasks the teams do not have time for — and leave without any guarantee of conversion, without formal recognition of their contribution, and sometimes without a usable reference on the job market. The infrastructure internship is structurally imbalanced: the company gets working capacity at marginal cost, the learner gets an experience that does not always translate into market access.

RATIO

◆ NASSIHA — THESE THREE PROFILES ARE NOT INTERCHANGEABLE

Their obstacles are distinct and their correction levers are too. Treating these three profiles as a homogeneous group would produce ill-suited recommendations. The IDPE proposal in section 5 primarily targets the self-taught. Apprenticeship is already a recommended entry lever — section 5 identifies what it does not solve. The intern question belongs more to employment law and conversion policy than to credential recognition.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE SECOND LOCK

CERTIFICATION AS AN ARTIFICIAL BARRIER TO ENTRY

To the diploma filter has been added a second lock, more recent and more insidious: vendor or publisher certification imposed as an implicit or explicit access condition on many infrastructure roles. AWS Solutions Architect, Azure Administrator, Cisco certification, VMware certification — these titles have become market signals that recruiters use like the diploma: a competence proxy, with the same advantages and the same flaws.

"Anatomy of the Loss" documented this mechanism from the angle of cognitive capture by hyperscalers: certification does not certify generic infrastructure competence — it certifies mastery of a specific product, in the version that was current when the exam was taken. It is a commercial retention tool as much as a competence recognition tool.

◆ THE CERTIFICATION PARADOX IN INFRASTRUCTURE

An AWS-certified candidate who has never seen a physical server room, never cabled a switch, never diagnosed a network fault in production at 3am is considered more qualified by many recruiters than a Blue with ten years of bare-metal field experience. This paradox is the direct symptom of the hard limits dissolution described in "Anatomy of the Loss": when physical competence disappears from evaluation criteria, what remains is mastery of service interfaces — and the certifications that validate them.

◆ THE INACCESSIBLE COST OF CERTIFICATION FOR THE SELF-TAUGHT

A cloud or vendor certification costs between €200 and €400 per exam, not counting preparation course costs. For a self-taught person without stable employment or an apprentice in training, this cost is prohibitive. The market therefore requires, to validate a competence often already acquired, a financial investment that excludes precisely the profiles who acquired that competence through their own effort rather than through a funded curriculum. It is an economic filter as much as a credential filter.

◆ THE FOUNDATION OF IRON AS A COUNTER-PROPOSAL

The pedagogical programme of this corpus — 26 weeks, from electricity to network security — was designed precisely to restore bare-metal competence, independently of any vendor certification. A learner who completes this pathway has a generic, transferable competence that no cloud certification guarantees. This is not an anti-certification stance — it is a reminder that foundational competence precedes and conditions the relevance of any certified specialisation.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

THE IDPE — A TOOL THAT EXISTS AND IS NOT USED

France has a remarkable legal specificity: the title of engineer is not regulated by law. Unlike other regulated professions (doctor, lawyer, architect), anyone can work as an engineer without a formal title — it is the market that decides who is an engineer. What is protected is the title of *graduate engineer*, awarded by schools accredited by the Commission des Titres d'Ingénieur (CTI).

Within this framework already exists an official mechanism, little known and greatly under-used: the **Ingénieur Diplômé par l'État (IDPE)** — the State-Certified Engineer. This state engineering title is accessible to working professionals who, without having followed a standard engineering curriculum, have demonstrated through their career and achievements a competence equivalent to the engineering level. The application is reviewed by a CTI-accredited engineering school, and the final decision belongs to an independent jury.

◆ THE CURRENT LOCK: THE ANNUAL JOURNAL OFFICIEL CALENDAR

The problem with the IDPE is not its design — that is sound. The problem is its access schedule: the opening of applications is announced in the Journal Officiel once a year, on an unpredictable date. For a professional who discovers this mechanism outside the annual window, the wait can last up to twelve months. For a self-taught person actively seeking employment, this administrative friction is often disqualifying. The Blue cannot wait for the next JO to be recognised.

◆ HOW THE IDPE CURRENTLY WORKS

1 — Application window announced in the Journal Officiel (once a year, variable date)
2 — Application submitted by the candidate to a CTI partner school
3 — Application reviewed by the school (analysis of career, achievements, competences)
4 — Appearance before a jury that rules on equivalence with the engineering level
5 — State-certified engineer title awarded if the jury is in favour
The process is rigorous and legitimate. The annual calendar constraint is the only artificial barrier to remove.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE PROPOSAL

PERMANENT IDPE — CONTINUOUS SESSIONS IN CTI SCHOOLS

The proposal is simple to state, precise to implement, and requires no new mechanism to be created: **REMOVE THE ANNUAL JOURNAL OFFICIEL CALENDAR CONSTRAINT AND ALLOW CTI-ACCREDITED ENGINEERING SCHOOLS TO PROCESS IDPE APPLICATIONS IN CONTINUOUS SESSIONS, YEAR-ROUND, ON A CASE-BY-CASE BASIS.**

The mechanism remains identical to today — application, school review, jury, decision. Only the time constraint is lifted. A Blue who discovers the IDPE in January can submit his application in January. A CTI school with a standing jury can review it in March. Recognition no longer depends on an unpredictable annual administrative window.

◆ WHAT THIS PROPOSAL CONCRETELY RESOLVES

A self-taught infrastructure professional with ten years of bare-metal production can have their level validated by a CTI school, obtain a state-certified engineer title, and present themselves to the market with official recognition of what they actually are. They stop being invisible to automated filters. They can apply for positions that formally require an engineering degree. The title reflects the competence rather than replacing it.

◆ OBJECTION 1 — "SCHOOLS HAVE NO INCENTIVE TO REVIEW THESE APPLICATIONS"

Processing an IDPE application represents a time cost for the school (file analysis, jury constitution and convening). This cost can be covered by processing fees, as is the case for VAE in other curricula. The school finds a complementary revenue model alongside its core activity, and a grounding in the local professional fabric. This is not a free service — it is a fee-based service that benefits both parties.

◆ OBJECTION 2 — "THIS WILL DEVALUE THE ENGINEERING TITLE"

The IDPE title is already awarded by a jury of graduate engineers, under the same standards as the CTI title. It is not a lightweight title — it is an alternative access pathway to the same title. Devaluation would come from lowering jury standards, not from opening the calendar. Making access continuous does not change the criteria — it removes an administrative friction that has no technical justification.

◆ LINK WITH VAE — A COHERENCE TO BUILD

The Validation des Acquis de l'Expérience (VAE — Recognition of Prior Learning) already allows degrees and titles to be obtained through recognition of professional experience. A permanent IDPE would be its specific version for the engineering title — with the additional rigour of the CTI jury. These two mechanisms should be explicitly articulated, to create a continuum of recognition for non-formal competences that runs from RNCP titles up to the state-certified engineer title.

◆◆◆

The Blue exists. He works in production. He maintains infrastructures that others cannot debug. He deserves a title that says what he is — not a blank page that says what he does not have.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

MANIFESTO · OPÉRATION DINDON · JUNE 2026

FROM THE WELL TO HEAD OF SRE

Eight Episodes · One Continuous Chain

Why the Corpus Is Not Theoretical

HUMAN

◆ **WHAT THIS DOCUMENT IS**

Opération Dindon is a corpus of more than thirty structural studies. Each thesis — the metal precedes the code, no sovereignty without matter, mastery of iron, transmission of tacit knowledge — has its origin in lived experience. This document traces that experience in eight episodes. It does not replace the studies — it says where they come from. And it answers the question every corpus reader eventually asks: who is the man writing this, and why should one believe him?

EPISODES	YEARS	WATERMARK
8	45+	HUMAN

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · *Opération Dindon* · *June 2026*

THE WELL · THE ELECTRICAL CABINET · THE SAKHR COMPUTER

◆ EPISODE 1 — THE HYBRID CLUSTER (0-6) · The real Bare-Metal

Amine is born into a world of contrasts. He grows up between town and primitive countryside — on one side the beginnings of modernity; on the other, a world without a tap or electrical socket. When water does not come from an API, you manage your own storage and latency: the path to the well.

This is the real Bare-Metal. Not as a datacentre concept — as daily reality. Resilience is not a property of a distributed system. It is a way of living when there is no safety net.

What this episode gives the corpus: "No Sovereignty Without Matter" is not an abstract thesis. It is a lesson learned before age six. Sovereignty means knowing where the water comes from — and being able to fetch it yourself when the API goes down.

◆ EPISODE 2 — THE SPARK AND THE SILICON (6-15) · The first mentor

A great-uncle, engineer at OCP Group, becomes the first maâllem. He opens electrical cabinets and shows how to tame lightning. Electrical current is not an abstraction — it is something one sees, hears, can touch (carefully) and can master.

Simultaneously, Amine gets his hands on a Sakhr computer. It is love at first sight — in both senses of the phrase. He understands that the world divides in two: those who submit to physics and those who programme it.

What this episode gives the corpus: the chain $V = R \times I$ → LED → motor → Arduino → GPU → AI finds its origin here. The OCP great-uncle is the first link in the educational chain that the corpus proposes to every physics teacher who reads the Anti-Amputation Foundation.

◆ CORPUS LINK — "DEAMPUTATION AT THE SOURCE"

What the great-uncle did — opening an electrical cabinet and showing the inside — is exactly what the Anti-Amputation Foundation asks secondary school physics teachers to do. The explicit invitation: "come and see what is inside." This gesture, repeated at national scale, is the deamputation. It worked for Amine. It can work for thousands of others.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE DISHES · THE SOLDERING IRON · SACRED RESPECT FOR THE METAL

◆ EPISODE 3 — THE AIRWAVE PIRATE (15-20) · The shadow antenna technician

Science lycée by day, shadow antenna technician by evening. Amine climbs rooftops to align satellite dishes. He dives into the electronics of satellite receivers for pocket money. He does not merely study automation and networking at school — he practises them in the field, soldering iron in hand.

Satellite signals are not abstractions in a physics lesson. They have a frequency, a polarisation, a measurable power level. A dish alignment to one tenth of a degree makes all the difference. Precision is not a quality — it is a necessity.

What this episode gives the corpus: "The Foundation of Iron" speaks of 26 weeks of training from electricity to Kubernetes. Episode 3 says that this foundation can also be learned on rooftops, at night, with a multimeter and personal motivation. On-the-job training is sometimes more effective than classroom training.

◆ EPISODE 4 — THE HARDWARE CRAFTSMAN (20-25) · Sacred respect for the metal

Before automating €18,000 servers, Amine works at the component level. Independent, he resurrects televisions, Nokia phones, liberates Canal+ decoders. This is where he forges what the corpus calls "sacred respect for the Metal".

For him, if you do not know what is under the bonnet, you control nothing. An engineer who has never opened a server, never replaced a power supply, never read BIOS logs understands infrastructure like a driver who does not know what is under the car bonnet. They manage in good weather. They are lost when it breaks down.

What this episode gives the corpus: "The metal precedes the code" is not a slogan. It is a conviction forged by resurrecting televisions in a workshop. And this conviction is the foundation of everything: the RSCNC32, the Moroccan Gibs, the Anti-Amputation Foundation, and the central argument of Opération Dindon.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

TELINF · EPF · ECRITEL · THE PROMISE KEPT EVERY NIGHT

◆ EPISODE 5 — MASTER OF THE MATRIX (25-32) · From Student to Teacher

Amine formalises his expertise at the Telinf centre in Paris. He is good enough to go from student to teacher. He teaches systems and networking while mentoring future engineers at EPF Engineering School in Sceaux. This is the era of transmission.

Transmission is not a complement to expertise — it is its ultimate test. An engineer who cannot explain what they do to someone who does not yet know it does not truly master their subject. Teaching forces clarity. It reveals what one believes one understands but does not yet understand well enough to explain.

Teaching since 2006 — this date in the signature is not a CV line. It is a continuity. Twenty years of transmission, first in an engineering school, now via the Opération Dindon corpus.

What this episode gives the corpus: "The Digital Craftsmen", "The Mentor and the Metal", "The Anti-Amputation Foundation" — these studies on transmission are written by someone who has taught. Not by someone who has read articles about teaching.

◆ EPISODE 6 — THE BAPTISM OF PRODUCTION (32-41) · The Sacred Fire of Operations

Theory is over — on to the sacred fire of operations at Ecritel. Amine starts in proximity support and ends as Head of Operations at OXYD. He manages datacentres, masters Linux servers, and learns that High Availability is not a concept — it is a promise kept every night on call.

This is where the experience documented in "The Accompanied Exodus" is born — migrations from A to Z, three-party meetings, root access, the TAD, the COPILs. Not technical literature — lived operational experience, night after night.

What this episode gives the corpus: "The Pager and the Body", "The Economy of Silence", "The Silence of the Rooms" — these studies on the invisible burden of operations are written by someone who has been on call. By someone who knows what it feels like to receive an alert at 3am and know they are the last resort.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

WEBORAMA · HEAD OF SRE · OPTIMISE THE IRON SO AS NOT TO FATTEN THE CLOUD

◆ EPISODE 7 — THE SRE AND EXTREME FINOPS (41-45) · The philosophy

Direction Weborama. Amine becomes SRE and Head of Datacentre. He manages massive data volumes and millisecond-precise response times. This is where he sharpens his philosophy: optimise the iron so as not to fatten the cloud unnecessarily.

FinOps of the Last Gram is not a technique — it is a posture. Every euro spent on compute must be justified by a real need, not by the convenience of ticking a box in a cloud console. Performance and sobriety are not antagonists — they are two sides of the same discipline.

This is also where the central conviction of Opération Dindon takes shape: hyperscaler cloud is not the future — it is one tool among others. And a tool whose physical substrate one does not master is a tool over which one is not sovereign.

What this episode gives the corpus: "FinOps of the Last Gram", "No Sovereignty Without Matter", "The Digital Ground" — these studies are written by someone who has optimised bare-metal servers against cloud offers, and made the real calculation, euro by euro, millisecond by millisecond.

◆ EPISODE 8 — HEAD OF SRE · The vision

The top of the pyramid. Head of SRE. Amine no longer merely manages servers — he defines the vision. He advocates GitOps, total automation and a return to pragmatism. He knows how to build complex infrastructures, but always keeps in mind the lesson of the well: simplicity and robustness above all.

The Head of SRE's vision is simple and direct: an infrastructure that works at 3am without human intervention is a well-designed infrastructure. An infrastructure that requires a permanent engineer on call is a poorly designed — or poorly documented — infrastructure.

What this episode gives the corpus: "Technical Primacy", "The Digital Craftsmen", "The Uniform of the Body" — these studies on technical authority and the visibility of the infrastructure body are written by someone who has exercised that authority, who has measured its absence, and who knows what is missing when it does not exist.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE CAUSAL CHAIN — EACH EPISODE IN THE CORPUS

WHERE EACH THESIS COMES FROM — AND WHY IT IS NOT THEORETICAL

EPISODE	EXPERIENCE	CORPUS STUDIES THAT FLOW FROM IT
1 · The Well	Life without API infrastructure	No Sovereignty Without Matter · The Digital Ground · أتستبدلون
2 · OCP + Sakhr	First maâllem, first machine	Anti-Amputation Foundation · Deamputation at the Source · The Mentor and the Metal
3 · Dishes	Field, precision, soldering iron	Foundation of Iron · Mastery of Iron · The Body That Makes · The RSCNC32
4 · Hardware	Component, respect for metal	The metal precedes the code · The Digital Iron · The Thought Machine
5 · Telinf / EPF	Student → Teacher since 2006	The Digital Craftsmen · The Replacement that Reveals · The Gentle Exit
6 · Ecritel / OXYD	Operations, on-call, datacentre	The Accompanied Exodus · The Pager and the Body · Economy of Silence · Departure of the Last One
7 · Weborama	SRE, FinOps, massive datacentre	FinOps of the Last Gram · Refurbished Grade A · MariaDB Without Trembling · Terms Under Microscope
8 · Head of SRE	Vision, GitOps, pragmatism	Technical Primacy · Uniform of the Body · Sovereign Interface · Global Anatomy of Amputation

From a well without a tap to the architecture of critical infrastructures.
From a soldering iron on rooftops to the definition of technical sovereignty.
From the OCP great-uncle who opens an electrical cabinet
to the thirty studies of Opération Dindon asking AWS to answer.
The chain is continuous. It has no rupture.
That is why the corpus is not theoretical.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

TEACHING GUIDE · OPÉRATION DINDON · JUNE 2026

THE MATHS TEACHER'S BAG

12 Objects · 12 Formulae · 90 Seconds

No Programme Change Required

HUMAN

◆ WHAT THIS DOCUMENT IS — AND WHAT IT IS NOT

This document requires **NO PROGRAMME CHANGE**. No extra session. No coordination with other teachers. No authorisation. No special training.

The teacher writes the formula on the board *exactly as they always have*. Then they place the object on the desk. 90 seconds. The lesson continues.

That is all. It is enough. And it is what some teachers already do — without naming it, without claiming it, without waiting for a reform to ask them to.

CARDS	DURATION	TOTAL BUDGET	WATERMARK
12	**90 SEC.**	**€20-30**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

Dedicated to all the teachers who answered differently.

HOW TO USE THIS DOCUMENT — THE GESTURE IN 3 STEPS

THE TEACHER WRITES · THE TEACHER PLACES · THE LESSON CONTINUES

1

THE TEACHER WRITES

They write the formula on the board as usual. Nothing changes in the lesson. The formula is there — sin(θ), πr^2, $a^2+b^2=c^2$.

2

THE TEACHER PLACES

They pull the object from the bag and place it on the desk — or pass it around. 90 seconds. They show the link between the formula and the object. One sentence is enough.

3

THE LESSON CONTINUES

The object stays on the desk or goes back in the bag. The lesson resumes exactly where it was. Nothing was sacrificed. A seed was planted.

◆ WHEN TO BRING OUT THE OBJECT — THE THREE NATURAL MOMENTS

Moment A — Just after writing the formula, before the first numerical example. The teacher writes the formula, turns around, pulls out the object: "Before we calculate — here is what this looks like in the real world." Then they put it away and resume.

Moment B — When a student asks "What is this good for?" The teacher pulls out the object: "Good question. Look." 90 seconds. Question answered. It is no longer hanging in the room.

Moment C — At the start of the lesson, before writing anything. The teacher places the object on the desk without a word. Students look. Then they write the formula. Curiosity is already there.

◆ HOW TO READ EACH CARD

🔢 **The formula**
Clearly visible. The starting point.

📦 **The object**
What to bring. Cost. Where to find it.

✋ **The steps**
What to do. 3 steps maximum.

❓ **The questions**
What to ask the class.

👷 **The trade**
Who uses this in real life.

⏱ **The moment**
When in the lesson.

HUMAN

◆ ONE RULE ONLY

The object does not replace the lesson. It does not supplement it. It gives it a physical meaning for 90 seconds. That is all. That is enough. The student who has seen the screw turn and the carriage advance will never again look at sin(θ) as an abstract symbol.

Amine RAITI · CC BY-NC-SA 4.0

01

ARITHMETIC · MULTIPLICATION

A × B = B × A

📦 THE OBJECT

Wooden lolly sticks × 100

€1 Supermarket / craft shop

✋ THE MANIPULATION

1. Arrange 3 rows of 4 sticks.
2. Flip: 4 rows of 3. Count the total.
3. Notice: same total — different arrangement.

❓ QUESTIONS FOR THE CLASS

Q1: Do 3 × 4 and 4 × 3 look the same physically?

Q2: Can you find an operation where the order of the gesture would change the physical result?

👷 IN REAL LIFE

BRICKLAYER / TILER

Laying tiles in 3 rows of 4 or 4 rows of 3 — same surface, same number, same cost.

Connection: Technology (grid layout) · Art (repeating patterns)

⏱ Moment in the lesson: Just after writing a × b = b × a on the board, before the first numerical example.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

02

ARITHMETIC · FRACTIONS

$\frac{1}{4} + \frac{2}{4} = \frac{3}{4}$

📦 THE OBJECT

An orange (or a chocolate bar)

€1 Supermarket

✋ THE MANIPULATION

1. Cut the orange into 4 equal quarters.
2. Place 1 quarter, then 2 quarters side by side.
3. Count: 3 quarters out of 4 — without calculating, just by looking.

❓ QUESTIONS FOR THE CLASS

Q1: How do you add ½ an orange and ⅓ an orange without calculating? What common division can you find?

Q2: Why can't you add ⅓ and ¼ directly without changing the division?

👷 IN REAL LIFE

COOK / PASTRY CHEF

Scaling a recipe from 4 to 6 servings: each ingredient × 6/4 = 3/2. Fractions are the tools of portioning.

Connection: Music (note values: crotchet = 1/4, minim = 1/2, semibreve = 1)

> ⏱ *Moment in the lesson: When a student asks why fractions with different denominators cannot be added directly.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

03

GEOMETRY · DEFINITION OF Π

π = C ÷ D

📦 THE OBJECT

String + 3 tins of different sizes

€0 Recycling — each student brings a tin from home

✋ THE MANIPULATION

1. Wrap string around a tin. Mark, unroll, measure: that is C.
2. Measure diameter D with a ruler.
3. Calculate C ÷ D. Repeat with all 3 tins. Compare results.

❓ QUESTIONS FOR THE CLASS

Q1: Why do we always get the same number regardless of the tin size?

Q2: Would it work with an oval object? Why not?

👷 IN REAL LIFE

BOILERMAKER / PLUMBER

Calculating the length of a pipe bend before cutting: length = π × diameter × (angle/360). Without π, the cut is wrong.

Connection: Physics (wheels, pulleys) · Technology (circular gears)

> ⏱ *Moment in the lesson: At lesson opening, before writing anything — place the 3 tins on the desk in silence.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

04

GEOMETRY · PYTHAGORAS

$A^2 + B^2 = C^2$

📦 THE OBJECT

Squared paper + scissors + dried beans

€0 School supplies + kitchen cupboard

✋ THE MANIPULATION

1. Cut 3 squares with sides of 3, 4 and 5 grid squares.
2. Fill the 3-square and 4-square with dried beans.
3. Pour all beans into the 5-square: it fills exactly.

❓ QUESTIONS FOR THE CLASS

Q1: Why do the two small squares fill the large one exactly?

Q2: Does it work with sides 5, 12 and 13? Calculate before checking.

👷 IN REAL LIFE

CARPENTER / TILER

The 3-4-5 rule: measure 3m on one wall, 4m on the other — if the diagonal is 5m, the angle is right. Every carpenter uses this without naming Pythagoras.

Connection: Construction (squareness check) · Navigation (GPS distance between two points)

> ⏱ *Moment in the lesson: Just after writing $a^2+b^2=c^2$ on the board — before any numerical example.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

05

TRIGONOMETRY · SINE AND COSINE

Y = R · SIN(θ)

📦 THE OBJECT

3 screws M6/M8/M10 20cm + nuts · cardboard crank-slider mechanism + paper fasteners

€3-5 AliExpress for screws · School supplies for crank

✋ THE MANIPULATION

1. Thread nut onto each screw. Turn 10 times — measure displacement. The ratio turns/distance is the pitch.
2. Compare the 3 screws: M6 (1mm pitch) → 10mm. M8 (1.25mm) → 12.5mm. M10 (1.5mm) → 15mm.
3. With the cardboard crank: turn the handle, watch the rod rise and fall. Measure position at 0°, 45°, 90°, 135°, 180°. Plot the curve: that is the sinusoid.

❓ QUESTIONS FOR THE CLASS

Q1: Why does the rod rise more slowly at the top than in the middle?

Q2: At what angle is the rod exactly at mid-height? Check with the formula.

👷 IN REAL LIFE

MECHANIC / ELECTRICIAN

AC mains is a sinusoid: $V(t) = 325 \cdot \sin(2\pi \cdot 50 \cdot t)$*. Each rotor revolution = one cycle. A piston displacement =* $R \cdot \sin(\theta)$.

Connection: Physics (oscillations, waves) · Music (pure tones = sinusoids)

> ⏱️ *Moment in the lesson: When the class sees sin(θ) for the first time — place the screws and crank before explaining the unit circle.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

06

TRIGONOMETRY · PARABOLA

$F = D^2 \div (16 \cdot d)$

📦 THE OBJECT

Small salvaged satellite dish + marble

€0-5 Car boot sale / recycling centre

✋ THE MANIPULATION

1. Place the dish open-face upward.
2. Release a marble from different points on the rim.
3. Observe: it always rolls to the same point — the focal point. That is where the LNB goes.

❓ QUESTIONS FOR THE CLASS

Q1: Why do all the marbles converge on the same point regardless of starting position?

Q2: What would happen with a square surface? With a hemispherical surface?

👷 IN REAL LIFE

SATELLITE TECHNICIAN / RADAR ENGINEER

Positioning the LNB at the exact focal point to receive the satellite signal. 1 cm off = signal lost. The formula is the precision tool.

Connection: Physics (parabolic mirrors, solar furnaces) · Astronomy (Cassegrain telescopes)

⏱ *Moment in the lesson: At the start of the quadratic curves chapter — before writing the equation $y = ax^2$.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

07 FUNCTIONS · PROPORTIONALITY
Y = K · X

📦 THE OBJECT

Kitchen scales + dried beans (50g minimum)

€10-15 Discount store / supermarket for the scales

✋ THE MANIPULATION

1. Weigh 10 beans, record. Weigh 20, 30, 50, 100. Plot the graph.
2. Notice: the line passes through the origin. That is $y = k \cdot x$. k = mass of one bean.
3. Add a fixed weight to the pan: the line no longer passes through the origin. That is $y = k \cdot x + b$.

❓ QUESTIONS FOR THE CLASS

Q1: How can you find the mass of a single bean without weighing it alone?

Q2: What does b represent physically when you add the fixed weight?

👷 IN REAL LIFE

PHARMACIST / ELECTRICIAN

Pharmacist: weight-proportional dosage — 5mg/kg × 60kg = 300mg.

Electrician: Ohm's law V = R·I — voltage is proportional to current. k = R.

Connection: Economics (fixed cost + variable cost) · Physics (Hooke's law)

⏱ Moment in the lesson: Just before introducing the gradient and y-intercept.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

08

ALGEBRA · EQUATIONS

2x + 3 = 11 → x = 4

📦 THE OBJECT

Kitchen scales + opaque bags + pebbles

€0 Recycling — garden pebbles, food bags

✋ THE MANIPULATION

1. Place 2 identical bags (unknown contents) + 3 visible pebbles on one side.
2. Place 11 pebbles on the other side. The scales balance.
3. Remove 3 pebbles from both sides. Divide by 2: each bag contains 4 pebbles.

❓ QUESTIONS FOR THE CLASS

Q1: Can you solve it without opening the bags? How?

Q2: What happens if you remove pebbles from one side only?

👷 IN REAL LIFE

ACCOUNTANT / SHOPKEEPER

I sold x items at €2 with €3 fixed costs and took in €11. How many items? The scales are elementary bookkeeping.

Connection: Chemistry (balancing a reaction) · Physics (equilibrium of forces)

> ⏱ *Moment in the lesson: When introducing the rule 'what you do to one side you do to the other' — the scales show why it is physical necessity, not arbitrary rule.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

09

STATISTICS · MEAN

$\bar{x} = \Sigma x_i / n$

📦 **THE OBJECT**

Tape measure — class heights

€2 DIY shop / already available

✋ **THE MANIPULATION**

1. Measure each student's height. Record on the board.
2. Calculate the sum, divide by the number of students.
3. Find the median: line everyone up by height, take the middle one.

❓ **QUESTIONS FOR THE CLASS**

Q1: Does the mean height correspond to anyone in the class?

Q2: Why is the median sometimes different from the mean?

👷 **IN REAL LIFE**

DOCTOR / NURSE

The growth chart compares a child's height to the average for their age group. Percentiles are the statistics of public health.

Connection: Biology (growth, BMI) · Economics (income distribution)

> ⏱ *Moment in the lesson: As an introduction to the statistics chapter — the data is real, students are invested.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

10

PROBABILITY

P(A) = FAVOURABLE CASES / TOTAL CASES

📦 THE OBJECT

6-sided dice + tally chart

€0 Board games already at school

✋ THE MANIPULATION

1. Roll the die 60 times. Record each result on the board.
2. Calculate the frequency of each face (count / 60).
3. Compare with theoretical probabilities (1/6 ≈ 0.167). Observe convergence.

❓ QUESTIONS FOR THE CLASS

Q1: After 5 consecutive sixes, has the probability of getting six on the next roll changed?

Q2: How many rolls does it take for frequencies to really resemble probabilities?

👷 IN REAL LIFE

INSURER / ACTUARY

Insurance premiums are calculated from the probability of a claim. An insurer who misjudges probabilities goes bust within 5 years.

Connection: Science (experimental protocol) · Games (poker strategy, backgammon)

⏱ Moment in the lesson: At lesson opening — roll the die before writing the definition of probability.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

11

SEQUENCES · EXPONENTIAL GROWTH

$U_n = U_0 \times R^n$

📦 THE OBJECT

One A4 sheet of paper

€0 Already in the classroom

✋ THE MANIPULATION

1. Fold in half, then in half again. Count layers at each fold (limit: 7 folds).
2. Calculate: initial thickness 0.1mm. After 20 folds → 100m. After 42 → Moon distance.
3. Note the physical impossibility of the 8th fold — the physical limit of a mathematical phenomenon.

❓ QUESTIONS FOR THE CLASS

Q1: From which fold would the stack exceed your height?

Q2: Why can a sheet of paper not be folded more than 7 times?

👷 IN REAL LIFE

BIOLOGIST / BANKER

Bacterial division: one bacterium → 2 → 4 → 8... In 24h at 20min/generation, one becomes 4 billion. Banker: compound interest — €1,000 at 5% over 40 years = €7,040.

Connection: Biology (cell division) · Computing (algorithmic complexity $O(2^n)$)

> ⏱ *Moment in the lesson: When introducing geometric sequences — fold the paper before writing the formula.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

12

SEQUENCES · FIBONACCI AND GOLDEN RATIO

$F_n = F_{n-1} + F_{n-2}$ · $\varphi \approx 1.618$

📦 THE OBJECT

An artichoke or a sunflower

€1 Market / grocery shop

✋ THE MANIPULATION

1. Count the spirals in the artichoke clockwise.
2. Count anticlockwise.
3. Notice: two consecutive Fibonacci numbers (8 and 13, or 13 and 21). Calculate the ratio.

❓ QUESTIONS FOR THE CLASS

Q1: Why does nature use this sequence to arrange its seeds?

Q2: Build a rectangle whose sides are in the ratio φ — does it look pleasing?

👷 IN REAL LIFE

ARCHITECT / DESIGNER / PHOTOGRAPHER

The proportions of a classical façade, the 16/9 screen ratio (≈ φ^2), the rule of thirds in photography — the golden ratio is the natural aesthetic of composition.

Connection: Art (composition, proportions) · Biology (phyllotaxis)

> ⏱ *Moment in the lesson: As an introduction to sequences — place the artichoke on the desk without a word. Let curiosity settle.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

SUMMARY TABLE · THE 12 OBJECTS AND THEIR BUDGET

ALL OF SECONDARY MATHS · FOR LESS THAN €30

N°	FORMULA	OBJECT	COST	TRADE
01	$a \times b = b \times a$	Lolly sticks × 100	€1	Bricklayer / Tiler
02	$\frac{1}{4} + \frac{2}{4} = \frac{3}{4}$	Orange / chocolate bar	€1	Cook / Pastry chef
03	$\pi = C \div D$	String + 3 tins	€0	Boilermaker / Plumber
04	$a^2 + b^2 = c^2$	Squared paper + dried beans	€0	Carpenter / Tiler
05	$y = R \cdot \sin(\theta)$	Screws M6/M8/M10 + crank	€3-5	Mechanic / Electrician
06	$f = D^2 \div (16 \cdot d)$	Satellite dish + marble	€0-5	Satellite tech / Radar engineer
07	$y = k \cdot x$	Kitchen scales + dried beans	€10-15	Pharmacist / Electrician
08	$2x + 3 = 11$	Scales + opaque bags + pebbles	€0	Accountant / Shopkeeper
09	$\bar{x} = \Sigma x_i / n$	Tape measure — class heights	€2	Doctor / Nurse
10	P(A) = favourable / total	6-sided dice	€0	Insurer / Actuary
11	$u_n = u_0 \times r^n$	A4 sheet of paper	€0	Biologist / Banker
12	$F_n = F_{n-1} + F_{n-2} \cdot \varphi \approx 1.618$	Artichoke or sunflower	€1	Architect / Designer

These 12 objects cover the entire secondary mathematics curriculum.
They fit in a bag. They cost less than €30 in total.
They each take 90 seconds.
And they finally answer the question every student has asked at least once:
"SIR, WHAT IS THIS GOOD FOR?"

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

MANIFESTO · OPÉRATION DINDON · JUNE 2026

SIR, WHAT IS THIS GOOD FOR?

The Question School Never Truly Answered

— and What It Cost

◆ **THE THESIS**

The contractual and architectural lock-in of large organisations is not a financial misstep. It is a cognitive bias cultivated over fifteen years of education that taught the future decision-maker that abstraction is superior to matter — and that touching the physical is a regression. This study traces the causal chain from the classroom to the boardroom, from sin/cos without application to the server rack one refuses to look at, from the question never truly answered to the cloud invoice never truly calculated. It is also a tribute to the mathematics teachers who answered differently — and who, without knowing it, changed trajectories.

CONDITIONING	VM COST / 5 YEARS	WATERMARK
15 YEARS	**€180,000**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

HUMAN

I

SECTION I · π WITHOUT A REAL CIRCLE — THE ARTIFICIAL RUPTURE

"SIR, WHAT IS THIS GOOD FOR?" — THE QUESTION SCHOOL NEVER TRULY ANSWERED

"Sir, what is this good for?"
— Every student, at some point in their schooling

◆ WHAT SCHOOL DOES — AND DOES NOT DO

School teaches π, sin, cos, If/Then structures, functions, vectors — on blank paper, in decontextualised exercises, with numerical values that represent nothing in the physical world. Mathematics is taught as an autonomous discipline with its own internal logic, its own notation, its own exercises. "Applied problems" exist — but as a complement, as a special case, never as a starting point.

What the student implicitly understands: mathematics matters in itself, independently of its application. The act of application — measuring, assembling, wiring, building — is secondary. And when they ask "Sir, what is this good for?", the answer is too often: "It's useful for higher education" or "It develops logical thinking" — true but abstract answers that show nothing, that do not let the student *see* what it is good for in their present life.

◆ WHAT THIS PRODUCES IN THE STUDENT'S MIND

A student who can calculate sin(30°) = 0.5 but does not know it is the ratio between two sides of a real triangle. Who knows π ≈ 3.14 but has never measured the circumference of a cylindrical object to verify it. Who solves quadratic equations but does not know they describe the trajectory of a projectile or the shape of a parabolic reflector.

When theory always precedes application — and when application never arrives — application becomes optional. And when application is optional, physical matter becomes foreign territory. Not incomprehensible — but unfamiliar, undesired, unvalued. The rupture between mind and matter is artificial. It is built by the pedagogical order. And it is durable.

HUMAN

◆ THIS IS NOT A CRITIQUE OF MATHEMATICS — IT IS A CRITIQUE OF THEIR ISOLATION

Mathematics is beautiful. π is a universal truth. Sines and cosines describe the world with a precision words cannot reach. This study does not say mathematics is bad. It says mathematics taught without ever showing what they are good for in the physical world produces decision-makers who have learned to value abstraction and to despise matter. And this contempt has a cost.

Amine RAITI · CC BY-NC-SA 4.0

SCHOOL BUILDS A HIERARCHY IT NEVER STATES EXPLICITLY — BUT EVERY STUDENT UNDERSTANDS

◆ THE IMPLICIT HIERARCHY

No teacher explicitly says: "Mathematics is superior to technology." Nobody writes in an official curriculum: "The manual gesture is inferior to abstract concept." But the student understands it anyway — through the implicit signals the system sends:

Students who succeed in mathematics go to science streams, then preparatory classes, then the grandes écoles. Students who struggle go to technical or vocational streams. The hierarchy of destinies mirrors the hierarchy of disciplines. Being "good at maths" opens all doors. Being "good with your hands" closes some.

The result: the act of transforming matter — the workshop, the wiring, the soldering iron, aligning a satellite dish — is associated with academic failure, not with intelligence. The student who succeeds in abstract mathematics learns to value abstraction. The student who struggles learns to despise what they are good at — because school told them it counts less.

◆ THE BIFURCATION AT AGE 12 — THE GENDER DIMENSION

"The Global Anatomy of Amputation" documented Stage -1 — the bifurcation at age 12 that specifically affects girls. The hierarchy of contempt applies differently by gender: the technical gesture is doubly marginalised for girls — too "manual" to be noble, and associated with a male world. A 12-year-old girl interested in antennae, circuits and machines finds no models in the cultural representations of her school environment.

It is not that girls are less capable of understanding sin and cos on a parabola. It is that nobody shows them a parabola. And when the parabola is not shown, neither is the fact that maths are what make it work. The hierarchy of contempt amputates twice: first the physical gesture, then girls from the physical gesture.

◆ WHAT THIS PRODUCES IN THE ADULT

The decision-maker who emerges from fifteen years of education with this hierarchy in their mind arrives in the boardroom with an unconscious bias: abstraction is intelligent, matter is thankless. Cloud — pure abstraction, invisible, no physical mess, sold with sophisticated technical vocabulary — is conformant with their value hierarchy. The server rack in the equipment room — physical, cabled, dusty, failing at 3am — is non-conformant. They do not need to do the financial calculation to choose. The school hierarchy chooses for them.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE DISH AND THE GUILLOCHÉS — LIVING PROOF

I ONLY UNDERSTOOD SIN AND COS ON A ROOFTOP · IN THE DARK · WITH A PROTRACTOR IN MY HANDS

◆ MOMENT 1 — THE SATELLITE DISH AND THE LNB

For years of mathematics lessons, sine and cosine were symbols on a page. We calculated them, memorised them, applied them in exercises. But there was no physical meaning — no reality to attach them to.

Then came the day of aligning a satellite dish. The question was concrete and urgent: where to place the LNB so that the satellite signal arrives correctly? The answer required calculating the focal distance based on the dish diameter and the satellite elevation angle. A protractor was needed. Sine and cosine were needed to determine the parabola's geometry. It was necessary to understand that the parabolic curve concentrates waves at a single point — the focal point — and that the LNB position must correspond exactly to this point.

At that moment, sine and cosine ceased to be symbols. They became tools on which the physical result depended: either the signal arrived, or it did not. No mark out of twenty. No correction possible the next day. Either it worked or it did not. And for it to work, one had to understand — truly understand — what these functions describe in real space.

The understanding gained that evening, on a rooftop, never left.

◆ MOMENT 2 — THE GUILLOCHÉS AND TRIGONOMETRY

A guilloché is a complex geometric pattern — rosette, spiral, interlace — whose main property is to make photocopying impossible. These patterns are produced by parametric equations involving trigonometric transformations: rotations, sinusoidal curves, spiral windings.

To create a security document with guillochés, it was necessary to understand what sin and cos do in two-dimensional space — how a trigonometric oscillation produces a curve, how the superposition of several oscillations produces a complex pattern not simply reproducible by photocopying. Trigonometry became necessary to produce something beautiful, useful, and technically irreducible to a simple copy.

Sin/cos/tan became the key to an aesthetic and a security property. Not an exercise.

HUMAN

◆ WHAT THESE TWO MOMENTS SAY ABOUT THE NATURAL ORDER OF UNDERSTANDING

In both cases, it was *necessity* that produced understanding — not academic constraint ("you must learn this for the exam") but real necessity ("if you do not understand this, your LNB will be mispositioned and you will have no signal"). Application created motivation. Motivation produced learning. And learning acquired in this context stayed — because it was anchored in a lived physical experience.

The natural order of human understanding is: **REAL PROBLEM → NECESSITY → LEARNING THE THEORY → DURABLE MASTERY**. School does the opposite: theory → exercises → optional application. And when application is optional, it disappears.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE BOARDROOM SWITCH — THE COGNITIVE BIAS IN ACTION

WHY THE CORPUS FIGURES ARE NOT ENOUGH — AND WHAT ELSE NEEDS TO BE UNDERSTOOD

The Opération Dindon corpus documented the 7.5× ratio between the GCP VM and bare-metal. It documented California jurisdiction in the Terms, noncancellable commits, egress fees, TSMC, the IME. These figures are true and verifiable. They are not enough to convince every CIO. This section explains why.

◆ IF LOCK-IN WERE PURELY ECONOMIC, THE FIGURES WOULD SUFFICE

A rational buyer comparing €180,000 vs €24,000 over 5 years chooses €24,000. They do not need 30 structural studies to be convinced. If CIOs continue choosing cloud after seeing the figures, their decision is not purely economic. It is also cultural. And cultural biases are not corrected by Excel spreadsheets.

◆ CLOUD IS CULTURALLY CONFORMANT · BARE-METAL IS CULTURALLY NON-CONFORMANT

Cloud-native conforms to what fifteen years of education taught one to value: modern, abstract, invisible, no physical mess, with sophisticated technical vocabulary (Serverless, Cloud-Native, Zero-Ops, Infinite Scalability). It resembles what school calls "intelligent".

Bare-metal does not conform to this hierarchy: physical, cabled, in a room that smells of circuit boards, failing at 3am and requiring one to put their hands in the rack. It resembles what school calls "manual".

The CIO does not choose cloud because they have not seen the figures. They choose it because cloud conforms to the value hierarchy that fifteen years of education built in their mind. And this hierarchy does not change with a financial argument — it changes with a foundational experience that shows physical gesture produces intelligence. Like the satellite dish on the rooftop.

◆ WHAT THE THOUGHT MACHINE ADDS

"The Thought Machine" documented that AI amplifies what the human brings. If the human was trained to bring only abstractions disconnected from reality — because school taught them that is all that matters — then AI amplifies disconnected abstractions. AI calibration is a physical skill as much as an intellectual one. Prompt engineering without grounding in understanding physical machines produces generic results. Amine RAITI's prompt engineering — grounded in twenty years of bare-metal operations — produces the Opération Dindon corpus.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE MATHEMATICS TEACHER AS ENTRY POINT

THEY HOLD THE KEYS TO ALL THE DOORS · THEY DO NOT ALWAYS KNOW IT

The central proposal of this study is not a curriculum reform. Not an additional budget. Not a new compulsory subject. It is a posture — that of a mathematics teacher who understands they are the gatekeeper of all other disciplines.

◆ WHY THE MATHS TEACHER IS THE ENTRY POINT

The mathematics teacher is the only one who holds the tools of all disciplines: sin and cos serve wave physics, servo technology, the art of guillochés and rosettes, computing algorithms for rotation, music frequencies. They teach before the others — their concepts arrive first in the student's journey. If they do not make the connection, other teachers receive students unable to see the continuity between theory and their discipline.

The idea is not that the maths teacher becomes a teacher of everything. It is that they pass the baton. They show the connection — and they say: *"To see how this applies in electronics, ask your technology teacher. To see how it produces a rosette, look with your art teacher. To see how it determines a satellite's position, explore with your physics teacher."* They are the conductor who gives the tempo — the others play their parts.

◆ THE CONCRETE PROPOSAL — ONE SESSION PER TERM

One session per term where the mathematics teacher brings a real-world object:

— A small satellite dish and protractor: "Calculate the focal length based on the diameter."

— A 3D-printed gear: "Find the gear ratio from the geometry."

— A security guilloché document: "Identify which trigonometric transformation produced this pattern."

— An architect's plan: "Calculate the real surface area from the scale and measurements."

— An Arduino and an LED: "Write the function that blinks the LED 440 times per second — the frequency of concert A."

Not in official curricula. Not assessed. Not graded. One hour per term during which the mathematics teacher tells their students: these symbols you are learning are not ends in themselves — they are tools. Here is what they are good for in the world.

HUMAN

◆ WHAT THIS PRODUCES ON THE SCALE OF A TRAJECTORY

A student who has seen, once, what sin and cos are good for in a satellite dish will never look at sin and cos the same way again. They will have an anchor. And this anchor, twenty years later, will be the difference between the CIO who signs for AWS because cloud conforms to their value hierarchy — and the one who says: "Wait, let us calculate the real TCO, and explain to me why we cannot have our own rack."

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE COMPLETE CAUSAL CHAIN

FROM YEAR 7 TO SIGNING FOR AWS — AND FROM THE DISH TO THE CORPUS

THE PUNITIVE ABSTRACTION CHAIN

Maths teacher: sin/cos without physical application

↓

Student: symbols without meaning in the real world

↓

Teenager: avoids technical streams (despised)

↓

Young adult: certifies in abstraction (AWS, GCP)

↓

Professional: values cloud over bare-metal

↓

CIO: signs for cloud without calculating real TCO

↓

Organisation: pays €180,000 over 5 years owning nothing

↓

Digital sovereignty: lost

THE COUNTER-CHAIN — WHAT THE DISH PRODUCED

OCP great-uncle: opens the electrical cabinet

↓

Physical curiosity awakened — real electricity

↓

Satellite dishes on rooftops — sin/cos become real

↓

Guillochés — trigonometry becomes aesthetic

↓

Workshop, components, soldering iron — metal respect

↓

Telinf, EPF — teaching, transmission

↓

Ecritel, OXYD — bare-metal operations, on-call

↓

Weborama — SRE, FinOps, iron optimisation

↓

Head of SRE → Opération Dindon

HUMAN

◆ WHAT THE TWO CHAINS SAY

The difference between the CIO who signs for AWS and the one who knows the corpus is not a difference of intelligence. Not a difference of academic training. It is a difference of informal education — a great-uncle who opens an electrical cabinet, a night on a rooftop with a protractor, a security document to guilloché.

Informal education should not be the only door to technical sovereignty. That is why the mathematics teacher is the entry point — because they have access to all students, not just those fortunate enough to have an engineer great-uncle.

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE PROPOSAL AND THE TRIBUTE

WHAT ONE HOUR PER TERM COULD HAVE CHANGED — AND WHAT SOME TEACHERS ALREADY DID

◆ THE TRIBUTE — TO TEACHERS WHO ANSWERED DIFFERENTLY

This study is not a charge against mathematics teachers. It is a tribute to the rare few who, at some point, pulled something out of their bag to show what maths are good for in the real world. These teachers exist. They are rare — not because they lack competence or curiosity, but because the institution does not ask them to do it. Pulling out a satellite dish in maths class means stepping outside official curricula. It is taking a minor but real professional risk. It is choosing to truly answer "Sir, what is this good for?" rather than responding "It's useful for higher education."

These teachers changed trajectories without knowing it. Some of their former students now manage critical infrastructure, train engineers, or write corpora on digital sovereignty. They do not know they planted the seed.

◆ THE PROPOSAL — WHAT IS IN THE CORPUS AND WHAT IS ADDED

The Anti-Amputation Foundation proposes 56 hours over 7 years to reverse the amputation. This study adds that the first gesture does not belong to the technology teacher — it belongs to the mathematics teacher. Before a student can understand why to blink an LED with an Arduino, they must understand that the mathematics they learn describe the physical world. And this connection, only the mathematics teacher can show first — because they arrive first in the student's journey.

Mastering tomorrow's great AI models requires minds capable of switching instantly from abstract equation to the muscular gesture of hardware configuration. This switching begins when the mathematics teacher pulls a satellite dish from their bag and asks the class: "*Who can tell me how to position the LNB to receive the signal?*"

Amine RAITI understood sine and cosine on a rooftop,
in the dark, with a protractor and an LNB in his hands.
This is not exceptional — it is normal.
What is exceptional is that school does not offer this to all students.
And this difference, thirty years later,
is measured in cloud invoices and lost sovereignty.

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE DIGITAL CRAFTSMEN

Elastic Syntec and Protection of the Infrastructure Body

From Notre-Dame to Datacentres

RATIO

◆ THE MISSING LAYER OF THE CORPUS (FRENCH CASE)

The Opération Dindon corpus built three layers: the Ticket and the Talent (nomenclature), the Foundation of Iron (training), the Uniform of the Body (visibility). The fourth was missing — legal protection. Without it, the other three remain recommendations without force. This study proposes the mechanism that makes them binding: elastic Syntec, applicable to infrastructure teams regardless of the sector employing them, on the model of the Compagnons du Devoir.

MODEL	TOOL	PROTECTION	WATERMARK
COMPAGNONS DU DEVOIR	**ELASTIC SYNTEC**	**ALL SECTORS**	**RATIO**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · NOTRE-DAME AND THE MISSING KNOWLEDGE

WITHOUT THE CRAFTSMEN, NOTRE-DAME WOULD LOOK LIKE LA DÉFENSE

On 15 April 2019, the spire of Notre-Dame de Paris collapsed. Hours after the fire, the question of reconstruction arose. Plans existed — archives, photogrammetric surveys, 3D models. What did not exist in sufficient quantity was the hand that knows how to transform these plans into physical reality. The knowledge of hewing oak as in the 13th century. The gesture of laying dressed stone without visible joints. The craftsman's bodily memory.

◆ THE COMPAGNONS DU DEVOIR — WHAT THEY PRESERVED

The Compagnons du Devoir is an organisation founded in the Middle Ages whose sole purpose is the transmission of tacit knowledge in manual trades. Carpenters, stonemasons, roofers, joiners. This knowledge is not in books — it is in hands, in gestures, in the daily decisions that only an experienced practitioner makes without realising it.

Without the Compagnons, Notre-Dame's reconstruction would have produced a technically correct and aesthetically approximate building. With them, the nave's timber frame was reconstructed in oak, using medieval techniques, with original assemblies — wooden pegs, no nails. This is not nostalgia: it is heritage sovereignty. France was able to rebuild Notre-Dame identically because an organisation had refused to let a knowledge die that nobody thought urgent to preserve.

◆ THE PARALLEL WITH DIGITAL INFRASTRUCTURE

When the last bare-metal SRE leaves — the one who can hear an anomaly in fan noise, read a network incident from raw metrics, who installed their first servers by hand before cloud existed — the runbooks will be there. The procedures will be documented. The AWS console will be available. But the tacit knowledge will have left with them. And the next major infrastructure crisis will produce an organisation that fumbles, calls consultants, starts from zero. Exactly like Notre-Dame without the Compagnons.

◆ NASSIHA — TACIT KNOWLEDGE DOES NOT WAIT

The Compagnons preserved their knowledge because they organised its transmission before the emergency arose. Not after the fire — before. The infrastructure body must do the same. "The Departure of the Last One Who Knows" is not a hypothetical future event — it is happening now, progressively, invisibly.

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE COMPAGNONS DU DEVOIR — ARCHITECTURE OF A SURVIVAL

WHAT THE INFRASTRUCTURE BODY CAN EXTRACT

The Compagnons du Devoir are not a living museum. They are 11,500 young people in training across 70 houses throughout France, supervised by experienced craftsmen, with an organisation structured around transmission and itineracy. Their architecture is transposable — not identically, but in its principles.

◆ THE FOUR PILLARS OF THE COMPAGNON ARCHITECTURE

The Tour de France: the young craftsman travels France — working successively with different masters, in different regions, exposing their knowledge to varied contexts. Not nomadism — structured training through contextual diversity. The Digital Craftsmen are open to all — equal access to grades regardless of gender is a founding condition of the association. The digital equivalent: structured rotation between different organisations (consulting firm, large corporation, public administration, startup, software publisher) for 12 to 24 months.

Compagnonnage: transmission does not happen in classrooms — it happens side by side, on the worksite, in real conditions. The master works, the apprentice observes, then does, then is evaluated. This is exactly the Observation Day documented in the corpus — extended to a structured relationship over months.

Grades: aspirant (in training) → compagnon (received) → master craftsman (transmitter). Each grade is open to all, equally. Each grade is earned through evaluation of the masterpiece — a concrete achievement judged by peers. The equivalent: aspirant SRE / compagnon SRE / master SRE, with a real technical project as reception masterpiece.

The Charter: the craftsman signs a body membership charter — commitment to transmission, quality, professional conduct. Not an employment contract — a commitment to the body.

◆ WHAT WE DO NOT COPY — AND WHY

The compagnon secret, initiation rites, esoteric symbols — these elements have their logic in the Compagnons' history but are not transposable to digital infrastructure. What we extract: the transmission structure, the grades, the formative itineracy, the membership charter. The form without the esoteric content.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE PROTECTION GAP — THE SRE ENGINEER IN A BANK IS NOT A BANKER

The Syntec collective agreement (IDCC 1486) is the reference convention for technical consulting firms, engineering companies and technology consulting firms in France. It covers approximately one million employees — including a large share of infrastructure engineers in consulting and technology firms. But it does not cover the SRE engineer working in-house at a bank, a hospital, an industrial company or a local authority.

◆ WHAT SYNTEC PROTECTS FOR THOSE IT COVERS

— Minimum salaries by coefficient (200 to 900+) with a three-category classification grid (ETAM, Engineers and Executives, Assimilated Executives).

— Regulation of day-rate arrangements (RTT, compensatory rest).

— Right to professional training with a specific budget.

— Enhanced healthcare and life insurance.

— Regulated mobility and relocation clauses.

— Redundancy payments above the legal minimum.

◆ THE PROTECTION GAP — INFRASTRUCTURE OUTSIDE THE TECH SECTOR

The SRE engineer in a bank is covered by the banking collective agreement (IDCC 2120) — a convention designed for banking roles, not technical ones. The SRE engineer in a hospital is covered by the hospital convention — designed for healthcare workers and administrators. The SRE engineer in the oil industry is covered by the petroleum convention. In all these cases, the applicable convention ignores the specificities of the infrastructure role: technical night on-calls, unmeasurable tacit knowledge, progression by competence rather than seniority, profile scarcity and the market salaries reflecting it.

◆ THE PRESS CONVENTION MODEL — PROOF THAT IT WORKS

Journalists benefit from the national journalists' collective agreement regardless of the sector of the company employing them — press, television, web, communications agency. A journalist in an advertising agency remains covered by the journalists' convention, not the advertising convention. This mechanism exists. It is legally established. It protects the profession wherever it is found, not wherever its employer is found. Infrastructure can and must work the same way.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · ELASTIC SYNTEC — THE PROPOSAL

THE LEGAL MECHANISM — PROTECTING THE BODY WHEREVER IT IS

Elastic Syntec is not a new collective agreement. It is an extension mechanism for the existing Syntec's scope of application, activated by three objective conditions, guaranteeing the infrastructure engineer the protection of their professional convention regardless of their employer's sector.

◆ THE MECHANISM — SECTORAL PRIMACY CLAUSE

Principle: when the employer's collective agreement is less protective than Syntec for the infrastructure roles listed in the nomenclature (SysOps/NetOps/OpInfra, six levels), Syntec applies by primacy. This mechanism already exists in French labour law — the favourability principle (Article L2254-1 of the Labour Code) allows a branch agreement to take precedence over a less favourable company agreement. Extension to inter-branch primacy requires a branch agreement or extension decree.

◆ THE THREE TRIGGER CONDITIONS

Condition 1 — Recognised role: the position is listed in the SysOps/NetOps/OpInfra nomenclature documented in "The Ticket and the Talent" — or in the official list of Syntec infrastructure roles to be created by branch agreement.

Condition 2 — Time threshold: the employee dedicates more than 50% of their effective working time to missions falling within the listed infrastructure roles. This threshold is documented by the job description and verifiable by HR audit.

Condition 3 — Title or certification: the employee holds an infrastructure RNCP title (Foundation of Iron pathway or recognised equivalent) or a professional certification recognised by infrastructure Syntec. This condition protects the mechanism against classification abuse.

RATIO

◆ WHAT ELASTIC SYNTEC TRIGGERS CONCRETELY

— Application of Syntec minimum salaries by coefficient, above the employer convention's minimums in under-paying sectors.
— Regulation of technical night on-calls (The Pager and the Body) with mandatory premiums.
— Right to infrastructure training (Foundation of Iron) on employer budget.
— Mobility clause protection — the infrastructure engineer cannot be relocated without respecting Syntec notice periods.
— Recognition of tacit knowledge in the classification grid — progression is not solely linked to seniority.

Amine RAITI · CC BY-NC-SA 4.0

REAL RIGHTS — NOT RECOMMENDATIONS

◆ PROPOSED SALARY GRID BY LEVEL

Aligned with the SysOps/NetOps/OpInfra nomenclature and existing Syntec coefficients:

Support / proximity technician (L1): coeff. 200-250 → €24,000-27,000 gross annual minimum.

Field technician (L2): coeff. 250-300 → €27,000-32,000 gross annual minimum.

Junior administrator (L3): coeff. 300-395 → €32,000-38,000 gross annual minimum.

SRE engineer (L4): coeff. 395-500 → €38,000-50,000 gross annual minimum.

Lead SysOps / Lead NetOps (L5): coeff. 500-660 → €50,000-65,000 gross annual minimum.

Principal SRE / Architect (L6): coeff. 660-900 → €65,000-90,000 gross annual minimum.

◆ ON-CALL — REGULATING THE PAGER

"The Pager and the Body" documented the cortisol debt of night on-calls. Elastic Syntec must specifically regulate infrastructure on-calls:

— Maximum 8 on-call weeks per quarter per engineer, mandatory documented rotation.

— 25% salary premium for any effective intervention between 10pm and 7am.

— Mandatory compensatory rest after any night intervention exceeding 2 hours.

— Prohibition of consecutive on-call exceeding 7 days without written employee agreement.

◆ TRAINING — THE FOUNDATION OF IRON AS A RIGHT

Syntec already requires a percentage of payroll dedicated to training. Elastic infrastructure Syntec mandatorily directs part of this budget toward recognised infrastructure training — Foundation of Iron, neutral certifications (Linux Foundation, HashiCorp, CNCF) — not exclusively hyperscaler cloud certifications that reinforce lock-in. This is the correction mechanism for the CPF Misused documented in the corpus.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE DIGITAL CRAFTSMEN — PROPOSED ARCHITECTURE

A PROFESSIONAL ASSOCIATION ON THE COMPAGNON MODEL

Elastic Syntec protects the body through law. The Digital Craftsmen preserve it through transmission. Both are necessary. Both are complementary. Syntec without transmission produces rights without knowledge. Transmission without Syntec produces knowledge without protection.

◆ PROPOSED STRUCTURE — THE DIGITAL CRAFTSMEN

Legal form: non-profit association, funded by member contributions and conventions with partner companies (which host craftsmen on rotation).

Grades: Aspirant SRE — any gender (on digital tour, < 3 years experience) / Compagnon SRE (received after masterpiece, 3-8 years) / Master Compagnon SRE (recognised transmitter, > 8 years, minimum mentoring commitment of 1 aspirant/year).

The Digital Tour of France: structured rotation of 12 to 24 months across 3 to 5 partner organisations from different sectors — a consulting firm, a major corporation, a public administration, a startup, a software publisher. Each rotation: 3 to 6 months minimum, with a Master Compagnon referent in place.

The Reception Masterpiece: real technical project carried out in operational conditions — infrastructure migration, deployment of a bare-metal monitoring system, construction of a CI/CD pipeline. Evaluated by a peer jury (Master Compagnons — all genders).

The Charter: commitment to transmission (minimum 1 mentored aspirant per Compagnon, any gender), technical quality, technological neutrality (no hyperscaler evangelism).

◆ THE LINK WITH THE TEACHER'S MANIFESTO

The Master Compagnon SRE — woman or man — is the corpus's Maâllem — the transmitter whose mechanism "The Mentor and the Metal" documented and "The Teacher's Manifesto" positioned. The Digital Craftsmen are fully open to women: equal access to all grades is a founding condition, not an aspiration. No longer an isolated individual transmitting in the shadows — a body that perpetuates itself, in all its diversity.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE MISSING LAYER — THE COMPLETE MODEL

FOUR LAYERS · ONE COMPLETE PROFESSIONAL BODY

The Opération Dindon corpus has built a four-layer model. Each layer is necessary. None is sufficient alone. Together, they constitute what no French organisation has yet built for its digital infrastructure — a complete professional body, named, trained, visible and protected.

◆ LAYER 1 — NOMENCLATURE
The Ticket and the Talent
SysOps / NetOps / OpInfra
6 levels · Non-Dilution Rule
→ Name the body

◆ LAYER 2 — TRAINING
The Foundation of Iron
26 weeks · FR and EN
Electricity to Kubernetes
→ Train the body

◆ LAYER 3 — VISIBILITY
The Uniform of the Body
5 levels · 5 signals
White coat · Black suit
→ Make the body visible

◆ LAYER 4 — PROTECTION
The Digital Craftsmen
Elastic Syntec + Compagnonnage
Rights wherever employed
→ **PROTECT THE BODY**

◆ INFRASTRUCTURE IS NOT A GOAL — IT IS THE FOUNDATION ON WHICH EVERYTHING RESTS

Digital infrastructure is not a sector like others — it is the substrate of all sectors. A bank without infrastructure stops. A hospital without infrastructure stops. A public administration without infrastructure stops. It is precisely because infrastructure is transversal that it must be protected by a transversal convention. Elastic Syntec says what French labour law has not yet said: the infrastructure engineer is not defined by the sector employing them — they are defined by the profession they practise. And this profession deserves protection commensurate with its actual importance.

◆◆◆

Without the Craftsmen, Notre-Dame would look like La Défense.
Without legal protection, infrastructure looks like any other employee.
These two realities share the same remedy: an organised, transmitting, protected body.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0
RAITI

TRAINING PLAN · 26 WEEKS · JUNE 2026

THE FOUNDATION OF IRON

Detailed Training Programme · 6 Months · 910 Hours

Restitution of the 2005 pathway, adapted for 2026 production readiness

26	910	10	35h
Weeks	Total hours	Pedagogical blocks	Per week

RATIO

◆ PROGRAMME OBJECTIVE

This plan details, week by week, the pedagogical foundation presented in the "Anatomy of the Loss" document of this corpus. The pathway follows a deliberate progression principle: three weeks of compressed conceptual fundamentals, then a constant climb toward directly employable skills — systems, middleware and networking, which represent 19 of the 26 weeks, or 73% of total time. The goal is not general technical culture: it is producing people ready for production, in six months, on a complete foundation from iron to network. Each brick in the final block strictly prepares the next, following a chain of technical dependency rather than an administrative category.

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

Operation Dindon

PROGRAMME OVERVIEW

Wk	Pedagogical block	Length	Theory/Practice
1-3	Electricity · Number conversion · Boole/Karnaugh	3 wk	55% / 45%
4-5	Physical automation (Arduino, PLCs)	2 wk	25% / 75%
6	Microcomputing (RAM, CPU, motherboard, BIOS)	1 wk	30% / 70%
7	Filesystems, disk images, dump/restore	1 wk	30% / 70%
8-16	**OS client/server on bare metal**	**9 wk**	28% / 72%
17-19	**Foundational networking** (OSI, addressing, VLAN, routing)	**3 wk**	32% / 68%
20	DHCP (on the networking base acquired)	1 wk	23% / 77%
21	DNS (direct prerequisite of Active Directory)	1 wk	23% / 77%
22-23	**Active Directory + GPO** (grouped together)	**2 wk**	27% / 73%
24	Web server	1 wk	23% / 77%
25	Database	1 wk	29% / 71%
26	**Network security (ACL) + Final synthesis**	**1 wk**	14% / 86%

◆ THE LOGIC OF PROGRESSION

The pathway follows a strict prerequisite chain: foundational networking (OSI, addressing, VLAN, routing) precedes every service that depends on it. DHCP is taught once subnets and VLANs are already mastered — addresses are distributed within a topology already understood. DNS precedes Active Directory, of which it is a genuine technical prerequisite. GPO is no longer taught in isolation midway through the OS block: it is grouped with Active Directory, immediately after, as the feature it actually is. Each week builds on the previous one's acquired knowledge rather than on an administrative category (OS / Middleware / Network) disconnected from the real technical logic.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

I

BLOCK I · WEEKS 1-3 · 105 HOURS

FUNDAMENTALS — ELECTRICITY, NUMBER CONVERSION, BOOLE/ KARNAUGH

W1 Electricity and power

50% T / 50% P

Theory (20h): voltage, current, power, Ohm's law, apparent/active/reactive power, kVA vs kW, reading nameplate data.

Practice (15h): multimeter measurements, calculating a server rack's load, simple power supply sizing.

Objective: read a UPS or PDU datasheet and calculate an admissible load.

W2 Number-base conversion

45% T / 55% P

Theory (15h): binary, octal, decimal, hexadecimal, arithmetic operations in each base.

Practice (20h): manual conversion exercises, calculating IP addresses in binary, subnet masks, hexadecimal colour conversion.

Objective: convert and calculate across all 4 bases without a calculator, on real networking cases.

W3 Boolean algebra, Karnaugh maps, logic circuits

35% T / 65% P

Theory (12h): AND/OR/NOT/NAND/NOR/XOR operators, truth tables, simplification via Boolean algebra and Karnaugh maps.

Practice (23h): designing a circuit from a specification (e.g. traffic light control), implementation on 74XXX-series integrated circuits, testing with LEDs/resistors/switches.

Objective: design and wire a simple logic circuit end to end.

◆ WHY THIS BLOCK IS COMPRESSED

These three weeks do not aim at expertise but at founding intuition: understanding that all infrastructure rests on electricity, that all data reduces to binary states, and that all computing logic is built from real physical gates. This is a conceptual foundation, not a professional end goal — it underpins the blocks that follow, which occupy the bulk of the training time.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

BLOCK 2 · WEEKS 4-5 · 70 HOURS

PHYSICAL AUTOMATION

W4 Embedded programming basics

25% T / 75% P

Theory (10h): microcontroller architecture, sensors and actuators, control loop concepts.
Practice (25h): first Arduino programs, sensor reading, motor/relay control.
Objective: make a program interact with a real physical event.

W5 Programmable logic controllers and integrating project

22% T / 78% P

Theory (8h): industrial PLC logic, automation specification.
Practice (27h): synthesis project combining sensors, conditional logic and actuators.
Objective: deliver a functional, documented mini automation project.

BLOCK 3 · WEEK 6 · 35 HOURS — MICROCOMPUTING

W6 PC and server hardware

30% T / 70% P

Theory (10h): components (CPU, RAM, motherboard, power supply, cooling), bus standards, BIOS/UEFI.
Practice (25h): full disassembly and reassembly of a workstation, hardware fault diagnosis, BIOS update, building a small server.
Objective: diagnose a hardware fault and assemble a machine end to end.

RATIO

BLOCK 4 · WEEK 7 · 35 HOURS — FILESYSTEMS AND DISK IMAGES

W7 Storage and data recovery

30% T / 70% P

Theory (10h): disk structure (sectors, partitions), filesystems (ext4, NTFS, ZFS), RAID concepts.
Practice (25h): creating partitions, taking disk images (dd, Clonezilla), restoring, recovering data from a corrupted disk.
Objective: back up, restore and recover a system from a disk image.

Amine RAITI · CC BY-NC-SA 4.0

5

BLOCK 5 · WEEKS 8-16 · 315 HOURS · 1/2

OS CLIENT/SERVER ON BARE METAL

The longest block of the programme — 9 weeks, 35% of total training time. This is the most directly marketable skill on the job market.

W8 Linux fundamentals (installation and shell)

29% T / 71% P

Theory (10h): Linux architecture, distributions, filesystem hierarchy.
Practice (25h): bare-metal installation, shell fundamentals, file management and permissions.
Objective: install and navigate a Linux system independently.

W9 Linux administration, level 1

23% T / 77% P

Theory (8h): package management, services (systemd), users and groups.
Practice (27h): package installation, service creation, permission management, system logs.
Objective: administer a basic Linux server in production.

W10 Linux administration, level 2

23% T / 77% P

Theory (8h): cron, shell scripts, log management, basic monitoring.
Practice (27h): automating routine tasks, maintenance scripts, setting up simple alerts.
Objective: automate routine administration tasks.

W11 Windows Server fundamentals

29% T / 71% P

Theory (10h): Windows Server architecture, roles and features, licensing.
Practice (25h): bare-metal installation, initial configuration, server role management.
Objective: install and configure a Windows server independently.

W12 Windows Server local administration

23% T / 77% P

Theory (8h): local users and groups, NTFS permissions, event logs.
Practice (27h): creating local accounts, managing file and folder permissions, troubleshooting via event logs.
Objective: administer a Windows server in standalone mode, without a directory.

◆ SEQUENCING NOTE

Group Policy Objects (GPO) have been removed from this week. GPO is an Active Directory feature, which is only introduced in week 22 — teaching it here would mean presenting a tool without the service it depends on. Week 12 therefore focuses on standalone Windows administration (local accounts, NTFS permissions), which logically prepares for directory-based administration, introduced later once networking and DNS have been acquired.

Amine RAITI · CC BY-NC-SA 4.0

OS CLIENT/SERVER — VIRTUALISATION, HA, SECURITY, SYNTHESIS

W13 Bare-metal virtualisation

29% T / 71% P

Theory (10h): type-1 vs type-2 hypervisors, virtual resource concepts, snapshots.
Practice (25h): installing a hypervisor (Proxmox VE or equivalent), creating and managing VMs.
Objective: deploy and manage virtual machines on a bare-metal hypervisor.

W14 High availability and backup

23% T / 77% P

Theory (8h): cluster concepts, RPO/RTO, backup strategies (GFS).
Practice (27h): configuring a simple cluster, setting up automated backups, restore testing.
Objective: ensure service continuity on a simple bare-metal environment.

W15 System hardening

29% T / 71% P

Theory (10h): OS hardening, update management, principle of least privilege.
Practice (25h): SSH hardening, local firewall configuration, configuration auditing.
Objective: secure a server according to basic best practices.

W16 OS synthesis project — Defence

14% T / 86% P

Theory (5h): case review, defence preparation.
Practice (30h): full deployment of a mini-environment (Linux server + Windows server + backup + hardening) to present at the defence.
Objective: demonstrate complete, autonomous systems administration competence.

◆ MID-PROGRAMME CHECKPOINT

At the end of this block, the trainee has installed, configured, secured and defended a complete bare-metal environment including two operating systems, virtualisation, backup and hardening. At this stage, this is already an employable skill for a junior systems technician role — the following two blocks (networking, directory and application services) add the depth needed for full autonomy.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

BLOCK 6 · WEEKS 17-19 · 105 HOURS

FOUNDATIONAL NETWORKING — OSI, ADDRESSING, VLAN, ROUTING

This block is now placed before all the network application services (DHCP, DNS, Active Directory) for which it is a direct prerequisite.

W17 OSI model and advanced addressing

34% T / 66% P

Theory (12h): the 7 OSI layers, encapsulation, IP addressing and masks, subnetting.
Practice (23h): addressing exercises, subnet calculations, packet capture (Wireshark).
Objective: understand and analyse a frame's journey through the OSI layers, master subnet calculation.

W18 Switching and VLANs

29% T / 71% P

Theory (10h): how a switch works, VLAN concepts, trunking.
Practice (25h): simulation on Cisco Packet Tracer or GNS3, VLAN configuration, network segmentation.
Objective: design and configure VLAN-based network segmentation.

W19 Routing

29% T / 71% P

Theory (10h): how routing works, routing tables, static vs dynamic routing (basics).
Practice (25h): configuring routers in a simulator, setting up inter-VLAN routing, connectivity testing.
Objective: configure functional routing between multiple VLAN-segmented networks.

RAFIO

◆ WHAT THIS BLOCK PREPARES

With OSI, addressing, VLANs and routing mastered, the trainee now has the complete topology on which the following three weeks will rely: DHCP distributes addresses within subnets that are now understood, DNS resolves names on infrastructure that is now mapped, and Active Directory will be installed on a network whose segmentation is already mastered.

Amine RAITI · CC BY-NC-SA 4.0

7

BLOCK 7 · WEEKS 20-21 · 70 HOURS

DHCP AND DNS — ON THE NETWORKING BASE ACQUIRED

W20 DHCP

23% T / 77% P

Theory (8h): how DHCP works, address leases, inter-VLAN DHCP relay.
Practice (27h): configuring a DHCP server, managing address pools across the VLANs already configured in week 18, testing across multiple subnets.
Objective: deploy a functional DHCP service on an already-mastered VLAN/routing topology.

W21 DNS

23% T / 77% P

Theory (8h): how DNS works, record types, recursive resolution.
Practice (27h): installing a DNS server, creating zones, resolution testing, DNS troubleshooting.
Objective: deploy and troubleshoot a DNS service — a direct technical prerequisite for the following week.

◆ WHY DNS IMMEDIATELY PRECEDES ACTIVE DIRECTORY

Active Directory technically depends on a functioning DNS service for locating domain controllers and resolving service records (SRV). Teaching DNS right before Active Directory is not an arbitrary scheduling choice: it is a genuine technical prerequisite, which prevents the trainee from configuring a directory on a service they do not yet master.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

8

BLOCK 8 · WEEKS 22-23 · 70 HOURS

ACTIVE DIRECTORY AND GPO — GROUPED TOGETHER

W22 Active Directory directory service / LDAP

29% T / 71% P

Theory (10h): directory concepts, LDAP structure, Active Directory principles, dependency on the DNS already deployed in week 21.
Practice (25h): installing a domain controller, creating users and groups, joining a client machine to the domain.

Objective: deploy a corporate directory and join client machines to it.

W23 Group Policy Objects (GPO)

25% T / 75% P

Theory (9h): GPO principles, inheritance and precedence, scope (site, domain, organisational unit).
Practice (26h): creating security and configuration GPOs, deployment on the machines joined to the domain in week 22, troubleshooting conflicting policies.

Objective: administer a Windows fleet via centralised group policies, on the directory deployed the previous week.

◆ SEQUENCING DESIGN

GPO is deliberately taught only after Active Directory, on which it technically depends, rather than earlier in the OS block. This grouping — Active Directory then GPO, across two consecutive weeks — respects the real dependency: one cannot administer group policies on a directory that does not yet exist.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

9

BLOCK 9 · WEEKS 24-25 · 70 HOURS

WEB SERVER AND DATABASE

These two bricks have a lighter networking dependency and can simply build on the base already acquired (addressing, DNS) without requiring any additional prerequisite.

W24 Web servers

23% T / 77% P

Theory (8h): HTTP/HTTPS fundamentals, client-server web architecture, SSL/TLS certificates.
Practice (27h): installing and configuring a web server (Apache/Nginx), hosting a simple site, resolved via the DNS already deployed, setting up HTTPS.

Objective: deploy a functional, secured web server, resolved by the infrastructure's DNS.

W25 Databases

29% T / 71% P

Theory (10h): relational model, basic SQL, database backup concepts.
Practice (25h): installing a DBMS (MySQL/PostgreSQL), creating databases, common queries, backup/restore, connecting from the previous week's web server.

Objective: install and administer a simple relational database, connected to a web application.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

10

BLOCK 10 · WEEK 26 · 35 HOURS

NETWORK SECURITY AND FINAL SYNTHESIS

W26 Network security (ACL) and final defence

14% T / 86% P

Theory (5h): firewalls, access control lists (ACL), security segmentation concepts, final defence preparation.
Practice (30h): configuring ACLs on the VLANs already segmented in week 18, designing and presenting a complete architecture integrating OS, networking, DHCP/DNS, Active Directory and application services deployed throughout the programme.
Objective: secure the entire infrastructure built through filtering, and demonstrate complete integration competence, from iron to network, ready for production.

◆ WHAT THE FINAL DEFENCE INTEGRATES

The week 26 defence does not cover a single isolated brick. It covers the entire chain built since week 8: Linux and Windows servers on bare metal, VLAN segmentation and routing, working DHCP and DNS, an Active Directory directory with group policies, web and database services, all secured by filtering rules. This is the demonstration that each week genuinely prepared the next, rather than a sequence of independent modules.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

CONCLUSION

WHAT THIS PROGRAMME PRODUCES IN SIX MONTHS

By the end of the 26 weeks, the trainee has demonstrated, through two defences (week 16 and week 26), complete integration competence: understanding the electricity and logic underlying all digital systems, automating a physical process, diagnosing and assembling hardware, managing storage at the disk level, installing and securing Linux and Windows operating systems in production, mastering foundational networking (OSI, addressing, VLAN, routing) before deploying the services that depend on it (DHCP, DNS, Active Directory and its group policies), then enterprise application services (web, database), and finally securing the whole through filtering.

◆ THE LINK TO THE REST OF THE CORPUS

This plan is not a theoretical proposal: it is the detailed restitution of the pathway completed by the author in 2005, at a vocational training centre for jobseekers, reordered here according to a strict chain of technical dependencies rather than by administrative category. Its relevance today directly responds to the cognitive mechanism described in "Anatomy of the Loss" — a complete technical framework, acquired in the right order, allows a professional to spontaneously propose architectures outside the trio cloud's sole ecosystem, because they have an end-to-end understanding of it, rather than a collection of disjointed modules.

This foundation does not oppose current DevOps or cloud skills. It complements them. A professional trained on this pathway, then later exposed to modern orchestration and automation tools, holds both frameworks — the abstraction and the hardware it covers — which is precisely the scarce competence this corpus identifies as the lever for reconquering digital sovereignty.

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE CPF MISUSED

Dormant rights, the Foundation of Iron, permanent IDPE

A funding architecture proposal

RATIO

◆ CONTEXT OF THE STUDY (FRENCH CASE)

This study is part of the Opération Dindon corpus, following "The Blue — What the Diploma Does Not Measure". It addresses the funding of infrastructure training and the recognition of non-formal competences. It puts forward a concrete proposal articulating three existing mechanisms — pooled CPF rights, France Travail, and VAE-IDPE funding under decree 2025-663 — to build a complete, profile-differentiated access route. No new legal mechanism is needed. The tools exist. The connection is missing.

Amine RAITI · Infrastructure Architect & SRE · Former instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · THE OBSERVATION

BILLIONS LYING DORMANT

The Compte Personnel de Formation (CPF — Personal Training Account) is credited at €500 per year for every employee, capped at €5,000 (€8,000 for less-qualified workers). Millions of accounts reach their ceiling without ever being used. The account holder does not know the rights exist, finds no relevant training in the catalogue, or deliberately preserves the balance for a future career change that never comes. The result is a vast stock of dormant rights — billions of euros immobilised in inactive accounts while the infrastructure sector deplores a shortage of trained profiles.

This paradox is structural. The CPF was designed to fund continuing professional training. In practice, it funds massively low-added-value short courses — driving licences, digital literacy certifications, office software qualifications — while the long, bare-metal training programmes that produce genuine systems and network technicians remain financially out of reach for the profiles who need them most.

◆ THE PROBLEM IS NOT THE AVAILABLE AMOUNT — IT IS THE ARCHITECTURE

€5,000 in a CPF account is enough to fund a significant portion of a quality infrastructure training programme. The Foundation of Iron — 26 weeks, 910 hours — is in the same cost range as the programmes funded by Pôle Emploi in the 2000-2015 period. The money is not the problem. The problem is that the money cannot find its way to the right training, for the right profile, at the right time.

◆ NASSIHA — A NOTE ON DORMANT RIGHTS

Dormant CPF accounts fall into three distinct categories: ignorance of the rights' existence, absence of relevant training in the catalogue, and deliberate saving for a planned future project. The third category must not be mobilised without explicit active consent. Any pooling mechanism must be strictly opt-in — never automatic.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE LESSON OF THE 2019-2022 FRAUD

The intuitive proposal is the simplest: allow holders of dormant CPF rights to transfer them to a colleague, a relative, or a training candidate who needs them. This proposal is legally prohibited — and the reason for this prohibition is documented by a recent episode.

Between 2019 and 2022, CPF fraud cost the Caisse des Dépôts several hundred million euros. Telephone canvassing networks usurped account holders' identities without their knowledge. Ghost training courses were billed and cashed without a single hour of training being delivered. The fraud mechanism relied precisely on the system's low friction: the less verification required, the easier the funds are to divert.

◆ WHAT FREE TRANSFER WOULD PRODUCE

A free transfer mechanism between natural persons instantly creates a secondary market for CPF rights. Intermediaries purchase dormant rights from holders who have no training intention, then resell access to organisations that bill non-delivered training. The logic is identical to the 2019-2022 fraud — free transfer is its natural vector.

◆ THE RIGHT ANSWER: CONTROLLED POOLING

Direct transfer is the wrong mechanism. Controlled pooling is the right one. A holder who explicitly and voluntarily decides to contribute to a dedicated paritaire fund — targeting training in shortage areas such as infrastructure and systems & networking — transfers their rights not to a person but to a mechanism managed by an accredited body. This is collective opt-in, not individual transfer. Minimal friction is maintained — it simply changes form: no longer administrative friction for the candidate, but explicit consent friction for the contributor.

◆ NASSIHA — WHAT SIMPLIFICATION MUST NOT MEAN

Simplifying access to training does not mean removing all friction. An active confirmation from the beneficiary remains necessary — it is the moment they validate their project and commitment. The friction to remove is the administrative burden of the application file, not the accountability of the candidate. A completely frictionless access system is a fraud waiting to happen.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THREE POPULATIONS, THREE MECHANISMS

A PROFILE-DIFFERENTIATED ARCHITECTURE

The most common design error in training funding reform proposals is to treat all profiles with the same mechanism. The funding needs of a jobseeker, an employed worker, and a self-taught person in career transition are structurally different. Their response must be too.

◆ POPULATION 1 — JOBSEEKERS AND CAREER-CHANGERS

Mechanism: France Travail — Conventioned Training Actions (AFC) and Operational Employment Preparation (POE).
Status: existing, proven, to be reactivated. The Foundation of Iron in its historical version was funded by Pôle Emploi under exactly this model. A Qualiopi-certified training provider delivering the Foundation of Iron can sign a direct convention with France Travail. No CPF required — no rights needed. France Travail funds the provider directly.
Advantage: covers precisely the profiles with few or no CPF rights — the self-taught Blue in career transition, the junior jobseeker, the intern without a stable contract.

◆ POPULATION 2 — EMPLOYED WORKERS WITH DORMANT CPF RIGHTS

Mechanism: opt-in CPF pooling fund managed by a paritaire body, dedicated to shortage-area training (infrastructure, systems & networking).
Status: to be created, via sector agreement or paritaire initiative. A holder voluntarily contributes to the fund. Rights are redistributed by the paritaire body towards eligible training, for beneficiaries identified by France Travail or OPCOs.
Advantage: mobilises the dormant stock without draining it anarchically — each contribution is a voluntary act, each redistribution is traced.

◆ POPULATION 3 — APPRENTICES AND UPSKILLING PROFESSIONALS

Mechanism: professionalisation contract + sector OPCO for apprentices; direct CPF + employer top-up for upskilling employees.
Status: existing. These mechanisms already work. The Foundation of Iron needs to be positioned as eligible training in the OPCO catalogues of the relevant sectors (digital, telecommunications, industry).

◆ NASSIHA — THE CPF IS NOT THE UNIVERSAL ANSWER

This proposal does not consist of "routing everything through the CPF". It consists of identifying the right mechanism for each profile. France Travail for the unemployed, pooled CPF for employees who voluntarily contribute, OPCO for apprentices. The CPF is one tool among three — not the sole solution.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE KEY FINDING

THE IDPE HAS BEEN FUNDABLE VIA CPF SINCE AUGUST 2025

The previous study in this corpus — "The Blue" — proposed a permanent IDPE accessible via continuous sessions in CTI schools. It also identified that the IDPE was not directly eligible for CPF funding. Further research reveals that the situation has evolved significantly since August 2025, and that a CPF funding pathway already exists for the IDPE — via the VAE (Recognition of Prior Learning).

◆ DECREE 2025-663 OF 18 JULY 2025: THE MISSING LINK

Decree n°2025-663, in force since 1 August 2025, extended CPF funding to the full VAE pathway — including accompaniment fees, preparation of the validation file, and jury fees. Before this decree, the CPF only partially covered accompaniment. Since 1 August 2025, a candidate pursuing the IDPE via VAE can fund their entire pathway through their CPF, provided the accompanying body is listed on Mon Compte Formation and Qualiopi-certified for the VAE category.

◆ THE FULL IDPE VIA VAE + CPF PATHWAY

Step 1: The candidate demonstrates 5 years of experience in engineering-level functions (without an engineering degree).
Step 2: They engage in a VAE pathway towards the IDPE, accompanied by an organisation listed on Mon Compte Formation.
Step 3: They mobilise their CPF to fund the accompaniment and jury fees (decree 2025-663).
Step 4: The CTI school reviews the file and convenes the jury.
Step 5: The jury rules — on success, the Ministry awards the state-certified engineer title.
Not covered by CPF: the initial admissibility examination fees (school administrative fees, variable) and jury attendance alone without accompaniment. These costs remain the candidate's responsibility or can be covered by the employer or OPCO.

◆ THE LAST LOCK: THE IDPE IS NOT ON THE FRANCE VAE PORTAL

The France VAE portal, created to centralise VAE processes, does not yet include the IDPE in its catalogue. The reason is technical and administrative: the IDPE is a state title awarded by the Ministry of Higher Education, outside the standard RNCP framework. A candidate searching for the IDPE on France VAE will not find it — they must go directly to Mon Compte Formation, which presupposes knowing that this pathway exists. This is not a legal lock. It is a visibility lock.

RAITIO

Amine RAITI · CC BY-NC-SA 4.0

THE PATH TO ELIGIBILITY

For the CPF to fund the Foundation of Iron directly as an initial training programme — not as a VAE — the training must be linked to a certification registered in the RNCP or the Répertoire Spécifique of France Compétences. This is a long process (12 to 24 months), rigorous, and costly. It cannot be improvised.

But there is a more immediate route than waiting for RNCP registration: the Foundation of Iron can be positioned as preparatory training for an existing RNCP-registered certification. Several level 5 or 6 certifications in systems and network administration are already registered. A provider delivering the Foundation of Iron can link their training to one of these existing certifications and gain immediate CPF eligibility — without waiting for a new RNCP registration.

◆ THREE CPF FUNDING ROUTES FOR THE FOUNDATION OF IRON

Route 1 — Linking to an existing RNCP certification: the training prepares candidates for an already-registered certification (Systems Administrator professional title, BTS SIO, or equivalent level 5-6). Immediate CPF eligibility. Timeline: a few weeks for listing on Mon Compte Formation.

Route 2 — Standalone RNCP registration: the Foundation of Iron becomes its own RNCP certification. Timeline: 12 to 24 months, with the cost of preparing the application and passing through the France Compétences commission. Full independence but significant upfront investment.

Route 3 — France Travail (AFC/POE): without an RNCP, immediately applicable for jobseekers via direct convention with France Travail. This is the historical route — it remains the fastest for the profiles with fewest CPF rights.

◆ THE RECOMMENDATION: START WITH ROUTES 1 AND 3 IN PARALLEL

Route 1 (existing RNCP link) provides immediate CPF access for employed workers. Route 3 (France Travail) provides immediate access for jobseekers. These two routes cover the bulk of the Foundation of Iron's target profiles from day one. Route 2 (standalone RNCP) is the 24-month objective — it grants full autonomy and maximum visibility on Mon Compte Formation.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE COMPLETE PROPOSAL

FOUR MEASURES, ZERO NEW MECHANISM

The entire proposal rests on legal and institutional mechanisms that already exist. No new legislation is needed. Four implementation measures suffice.

◆ MEASURE 1 — LIST THE IDPE ON THE FRANCE VAE PORTAL

A DGESIP (Directorate-General for Higher Education) instruction is sufficient to list the IDPE on the France VAE portal. This measure makes the CPF→VAE→IDPE pathway visible and accessible without the candidate having to discover it themselves. It is two lines in a ministerial circular. It connects an existing funding mechanism (CPF, decree 2025-663) to an existing title (IDPE) via an existing portal (France VAE).

◆ MEASURE 2 — MAKE THE IDPE PERMANENT VIA CONTINUOUS CTI SESSIONS

Remove the annual Journal Officiel calendar constraint and allow accredited CTI schools to process IDPE applications year-round, in continuous sessions. Same process, same jury, same standards. This measure is inseparable from Measure 1: an IDPE visible on France VAE but accessible only once a year remains unusable in practice.

◆ MEASURE 3 — CREATE AN OPT-IN CPF POOLING FUND FOR INFRASTRUCTURE

Via a sector agreement in the digital and telecommunications industries, create a fund fed by voluntary contributions from CPF holders with dormant rights. This fund finances exclusively shortage-area training identified by the sectors — infrastructure, systems & networking, SRE. It is managed by an accredited paritaire body with full traceability of all flows.

◆ MEASURE 4 — CONVENTION THE FOUNDATION OF IRON WITH FRANCE TRAVAIL

A provider delivering the Foundation of Iron obtains Qualiopi accreditation and signs an AFC convention with France Travail. Funding for jobseekers and career-changers is assured without CPF. This is the historical route — it worked for fifteen years. It does not need to be reinvented. It needs to be reactivated.

◆◆◆

The money is there. The mechanisms are there. The title is there. What is missing is the connection between them — and that connection requires four implementation measures, not a legislative reform.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART VIII

SOLUTIONS: EXIT

13 pieces

HUMAN

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

TECHNICAL PRIMACY

The Principal SRE's Final Word

Budget Oversight vs Technical Decision Authority

RATIO

◆ THE CENTRAL CONCEPT

Technical Primacy is the constitutive right of the Principal SRE to have the final word on architecture and operational decisions, within the budget set by management. It does not contest managerial authority — it delimits it. Management sets the envelope. The Principal SRE decides how to spend it and which technical risks to refuse. Confusing these two roles systematically produces Invisible Debt, the Infrational Loop, and the Departure of the Last One Who Knows.

CONCEPT	MODEL	WATERMARK
TECHNICAL PRIMACY	**CHIEF PHYSICIAN**	**RATIO**

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

THE CONFUSION THAT COSTS

In most organisations, nobody has explicitly defined who has the final word on infrastructure decisions. There is a budget, a committee, an IT manager, sometimes a CIO. And in arbitration situations — migrate or maintain, buy or build, accept or refuse a risk — the final decision goes to whoever holds the highest title in the room. Not to whoever knows the subject best.

◆ THE TWO ROLES — THE MISSING DISTINCTION

Budget oversight belongs to management. Setting the annual envelope, arbitrating between competing projects, prioritising according to the organisation's strategy. This right is legitimate, necessary, and requires no technical background. A CEO or CFO can decide that infrastructure gets €800,000 this year and not €1.2 million. That is a management decision.

Technical decision authority belongs to the Principal SRE. Within that €800,000: which architecture, which tools, which suppliers, which risks accepted or refused, which technical debt taken on or repaid. This decision requires expertise that general management does not possess — and cannot possess without having practised the profession.

Confusing the two roles systematically produces the same result: technically suboptimal decisions made by people who do not measure their consequences, imposed on teams that can see them coming and cannot stop them. The debt accumulates. The incident eventually happens. And the Principal SRE responds at 3am for a decision they did not make.

◆ THE WARNING SIGNAL — WHEN THE CONFUSION IS IN PLACE

Three indicators that Technical Primacy is absent in an organisation:

— Infrastructure arbitrations happen in executive committees without the Principal SRE present, or without their opinion being constitutive of the final decision.

— The Principal SRE says "I advise against this migration now" and the migration launches anyway without formal documentation of the accepted risk.

— Post-decision incidents are analysed as technical failures rather than as consequences of uninformed managerial decisions.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE CHIEF PHYSICIAN AS MODEL

THE HOSPITAL SOLVED THIS PROBLEM — INFRASTRUCTURE CAN DO THE SAME

The hospital is the organisation that has best resolved the tension between managerial authority and technical expertise. The solution has been institutionalised for a century and a half. It is so obvious that nobody questions it. It is called the separation between administrative management and medical authority.

◆ THE HOSPITAL MODEL — INSTITUTIONALISED SEPARATION

The hospital director sets the cardiology department's budget. They decide the envelope, staffing levels, equipment investments. They do not decide the treatment protocol for a patient. If they decided that a patient receives a particular treatment for budgetary reasons against the head physician's advice, this would constitute endangering human life — in the legal sense. The department's chief physician has the final word on medical decisions, within the budget set by management.

This separation is so obvious in the medical world that stating it seems trivial. It is not trivial — it is the result of decades of institutional construction following real catastrophes. Digital infrastructure is going through the same catastrophes, with the same causes, without having yet built the same protection.

◆ THE INFRASTRUCTURE TRANSPOSITION — TERM BY TERM

Hospital director → CEO / CIO / CFO: sets the infrastructure budget, arbitrates between projects, prioritises according to strategy. Full oversight of the envelope and priorities. No decision rights on technical choices within that envelope.

Department chief physician → Principal SRE: within the set budget, decides on architecture, tools, suppliers, risks accepted or refused. Their opinion is constitutive and documented. Their refusal of a risk is formal and recorded.

Endangering the patient → Major infrastructure incident: when the director imposes a medical protocol, they endanger a life. When management imposes a technical decision, they create invisible debt that materialises as an incident. The difference: the infrastructure incident arrives 18 months later, not the next day. This delay masks causality — and allows the confusion to perpetuate.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · TECHNICAL PRIMACY — DEFINITION AND SCOPE

WHAT IT COVERS · WHAT IT DOES NOT COVER

Technical Primacy is not a blank cheque. It is not the Principal SRE's right to decide without accountability, nor to indefinitely block strategic decisions. It is a precise, delimited, defensible perimeter — that protects the organisation as much as the Principal SRE.

◆ WHAT TECHNICAL PRIMACY COVERS

Architecture: infrastructure pattern choices, decoupling or integration decisions, choices between bare-metal and cloud approaches, redundancy level definitions.

Operations: incident procedures, maintenance windows, alert thresholds, production rollback decisions.

Technical security: refusal of a configuration deemed insecure, even if requested by a business project or management. This refusal is documented, motivated, and formally recorded.

Documented risk refusal: the Principal SRE can refuse to sign off on a production deployment or migration if they judge the risk unacceptable. This refusal is formal. Management may override — but formally assumes responsibility for the risk by signature.

◆ WHAT TECHNICAL PRIMACY DOES NOT COVER

The overall budget: the Principal SRE does not set their own budget. They propose, argue, but the final decision belongs to management.

Strategic priorities: which infrastructure project comes first — if the decision is strategic rather than purely technical, management arbitrates.

Headcount: how many people in the infrastructure team — an HR and budgetary decision, not Technical Primacy.

Overall project schedule: when an infrastructure project must be delivered — within technically reasonable limits, the schedule is a managerial decision.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THREE SCENARIOS DOCUMENTED BY THE CORPUS

◆ SCENARIO 1 — INVISIBLE DEBT (Corpus: The Invisible Debt)

Management decides on a rushed cloud migration to meet a contractual deadline. The Principal SRE indicates that the architecture is not ready, that the migration will create uncontrolled dependencies and an estimated 18 months of corrective technical debt. Their recommendation is ignored. The migration happens on schedule. Eighteen months later, the infrastructure team spends 60% of its time repaying the debt — less time for new initiatives, less time for training, less time for knowledge transmission. Management wonders why infrastructure is "slow". Causality is invisible because the delay between decision and consequence is 18 months.

◆ SCENARIO 2 — THE INFRATIONAL LOOP (Corpus: The Infrational Loop)

The Principal SRE recommends investing in bare-metal training for the team — rare skills, essential to the organisation's autonomy. Management decides to prioritise hyperscaler cloud certifications — cheaper, more visible on CVs, better for team "visibility". Two years later, the organisation is more dependent on hyperscalers than before, engineers are less autonomous on bare-metal, and the Principal SRE — whose recommendations have been systematically ignored — starts looking at market offers.

◆ SCENARIO 3 — THE DEPARTURE OF THE LAST ONE WHO KNOWS (Corpus: The Departure)

The Principal SRE has refused a technical decision three times in two years that they judged risky. All three times, management overrode without formally documenting the risk assumption. Two of the three incidents happened as predicted. After the third, management asked the Principal SRE "why doesn't infrastructure deliver on its commitments". They leave. Ten years of accumulated tacit knowledge leave with them. The organisation hires a consultant at €1,200/day to understand the existing system. It will never recover the level of mastery it had.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

WHAT MANAGEMENT DECIDES — AND WHAT IT PROTECTS BY NOT DECIDING

This study is not a critique of management — it is a protection of management. A CEO who makes a bad technical decision because nobody has clearly told them "this is not your decision to make" is an exposed CEO. Technical Primacy protects them as much as the Principal SRE — it clearly states where their responsibility ends.

◆ WHAT NON-TECHNICAL MANAGEMENT LEGITIMATELY DECIDES

The budget envelope: how much the organisation invests in its infrastructure. This is a strategic and financial decision — it belongs to management.

Priorities between projects: if the organisation must choose between migrating the datacentre and building a new data platform, management arbitrates according to strategy. They inform the Principal SRE of the priority — the Principal SRE says how to implement it technically.

The acceptable risk level: management can decide that the organisation accepts a risk the Principal SRE advises against — but this must be an explicit, documented, signed decision. Not a silence that masks risk-taking.

The strategic schedule: when a system must be delivered, according to contractual and strategic constraints. Within this, the Principal SRE states what is technically feasible or not.

◆ NASSIHA — THE MANAGEMENT THAT PROTECTS ITS ORGANISATION

The management that recognises the Principal SRE's Technical Primacy does not abandon its authority — it exercises it correctly. It says: "I decide the envelope and the strategy. For technical decisions within that envelope, I trust the Principal SRE and hold them accountable." This is the most effective delegation contract that exists in a technical organisation. And when something goes wrong despite Technical Primacy, responsibility is clearly located — not diluted in a poorly documented collective decision.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · TECHNICAL PRIMACY IN THE CORPUS

WHAT THE CORPUS ALWAYS SAID WITHOUT NAMING IT

Technical Primacy is the concept that makes the entire Opération Dindon corpus coherent. Every study documented a consequence of its absence. None had yet named the absence itself. This concept is the piece that closes the puzzle.

◆ NO SOVEREIGNTY WITHOUT MATTER — THE FOUNDING THESIS

The founding thesis says: the human is sovereign when they can act on the matter they depend on. Applied to the organisation: an organisation is sovereign over its infrastructure when the Principal SRE can decide on that infrastructure. Without Technical Primacy, the organisation delegates its technical sovereignty to people who do not understand it. This is extraterritorial cloud at the internal scale — not a foreign hyperscaler deciding, but non-technical management deciding what they do not understand.

◆ THE UNIFORM OF THE BODY — THE SUIT WITHOUT AUTHORITY IS DECORATIVE

The Lead SRE's black suit and red tie signals to the organisation: this person is a technical commander. If this commander does not have the final word on technical decisions, the suit is an empty sign — it signals an authority that does not exist. Technical Primacy is the condition that gives the uniform its substance. Without it, one dresses an adviser. With it, one recognises a commander.

◆ THE DIGITAL CRAFTSMEN — THE GRADE WITHOUT POWER IS HONORARY

The Master Compagnon SRE grade only makes sense if the master level actually gives the final word on technical decisions in their organisation. If the Master Compagnon SRE is systematically overridden by non-technical management, their grade is honorary — it rewards knowledge that has no power. Technical Primacy is the condition that transforms the grade into real authority.

◆ ELASTIC SYNTEC — PROTECTION WITHOUT DECISION IS HOLLOW

Elastic Syntec protects the infrastructure body's rights — salaries, on-calls, training. But a protected body that does not have the final word on its domain of expertise is a body paid to execute decisions it did not make. Technical Primacy is the layer that completes protection — protecting not only salary and conditions, but the right to decide.

Amine RAITI · CC BY-NC-SA 4.0

FOUR CONCRETE MECHANISMS

◆ MECHANISM 1 — THE PRIMACY CLAUSE IN THE CONTRACT

The Principal SRE's contract includes an explicit Technical Primacy clause — defining their autonomous decision perimeter, their documented veto right, and the formal escalation procedure when management wishes to override. This clause does not give the Principal SRE absolute power — it clearly delimits responsibilities. It protects both parties.

◆ MECHANISM 2 — THE DOCUMENTED VETO

When the Principal SRE opposes a technical decision, their opposition is formal and written. The veto document specifies: the contested decision, the identified technical risks, the estimated debt created or probable incident, and the alternative recommendation. This document is retained. If management decides to override, they sign the document — formally assuming responsibility for the risk. This mechanism protects the Principal SRE from responsibility for incidents they did not cause, and protects management against undocumented decisions.

◆ MECHANISM 3 — PRESENCE IN THE EXECUTIVE COMMITTEE

The Principal SRE is present in all committees where decisions with infrastructure impact are made — not as an occasional invitee, but as a permanent member with deliberative voice on technical subjects. This presence is structural, not circumstantial. It ensures arbitrations happen with correct technical information, not after the fact.

◆ MECHANISM 4 — FORMAL ESCALATION

When management wishes to override the Principal SRE's veto, a formal escalation procedure is triggered: consultation of an external technical third party (solutions architect, infrastructure audit firm), formal meeting with documented arguments from both parties, final decision signed by the appropriate management level with explicit mention of risk assumption. This procedure is not an obstacle — it is a guarantee that the decision is made with full knowledge.

RATIO

The Principal SRE is not infallible.
But they are the only one who measures what they decide.
Giving them the final word gives the organisation
its best chance of not regretting its technical decisions.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE GENTLE EXIT

A Proposed Conversion Without Rupture

For DevOps, Cloud Architects and Organisations in Transition

◆ THE POSTURE OF THIS STUDY

This study does not say "you were wrong". It says "here is what is missing and how to add it — with today's tools". The Opération Dindon corpus diagnosed the problems. This study proposes exit paths. The SysOps/NetOps/OpInfra nomenclature proposed in "The Ticket and the Talent" is not an adopted framework — it is a pragmatic proposal. A plan to exit collapse, not a statement of the current state. It also opens a door nobody has yet named: the entry of women into the infrastructure body.

STATUS	APPROACH	WATERMARK
CORPUS PROPOSAL	**TODAY'S TOOLS**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · NOT A FAULT — A GAP

THE DEVOPS WAVE WAS A REAL RESPONSE TO A REAL NEED

People who carry the DevOps title today are not enemies of the infrastructure body. Many were poorly oriented, poorly trained, poorly named — not through bad faith, but because the market called anything and everything "DevOps" for ten years. They have the foundations. A layer is missing — not all the competence. The gentle exit does not say "you were wrong" — it says "here is what is missing and how to add it, with what exists today".

◆ WHAT THE DEVOPS WAVE ACTUALLY BROUGHT — ASSETS TO PRESERVE

Three real and valuable things: collaboration between developers and operators, deployment automation (IaC, CI/CD), and a culture of measurement (monitoring, alerting, SLO, post-mortem). These are assets. The gentle exit anchors them in the physical reality they were missing — it does not erase them.

The problem is not that DevOps engineers exist. It is that the title progressively replaced precise technical titles without transmitting their fundamental competences: hardware, physical networking, the low-level OSI stack, storage, bare-metal. Competences in height, fragile in depth.

◆ THE NOMENCLATURE — A PROPOSAL, NOT AN ESTABLISHED FACT

The SysOps/NetOps/OpInfra nomenclature proposed in "The Ticket and the Talent" is not a framework adopted by any professional branch or institution. It is an Opération Dindon corpus proposal — a plan to exit collapse, documented, argued, operational. It proposes six levels, three categories, defined competence criteria. It only exists if organisations choose to adopt it. This study proposes the tools to make that adoption possible — progressively, without social rupture.

◆ THE THREE POPULATIONS THAT NEED AN EXIT

"Orphan" DevOps: real competences, a title that no longer serves them. They sense the gap — they do not know how to fill it without starting from zero.

Cloud architects without roots: legitimacy on cloud certifications, fragility below. They want to anchor — not rebuild everything.

Organisations in transition: teams built on the DevOps wave, corpus problems materialising — debt, departures, dependence. They want to rebuild without social rupture or public admission of error.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE PROPOSED CONVERSION MAP

IF THE NOMENCLATURE WERE ADOPTED — WHAT EACH DEVOPS WOULD GAIN

If an organisation decided to adopt the nomenclature proposed in "The Ticket and the Talent", here is how the conversion could work — without downgrade, without rupture, valorising existing competences. This table is a proposed reading, not a decision already made.

CURRENT TITLE	PROPOSED TITLE	GAP · SOLUTION AVAILABLE TODAY
Junior DevOps / Junior Ops	**SysOps L2 / L3** (proposed)	Hardware, physical network, storage · Foundation of Iron modules 1-12 · LFCA Linux Foundation (€395)
Senior DevOps / Platform Engineer	**SysOps / NetOps L4** (proposed)	Physical OSI, routing, bare-metal · LFCS (€395) + shadowing Principal SRE
DevOps Lead / SRE Lead	**Lead SysOps L5** (proposed)	Major bare-metal incident, multi-site architecture · CKA CNCF (€395) + 6-month rotation
Cloud Architect / Solutions Architect	**Principal SRE L6** (proposed)	Physical substrate, Technical Primacy · Observation Day + Foundation of Iron 1-8 + mentoring
DevOps Coach / Agile Infra Coach	**Senior SRE Transmitter** (proposed)	Formalised mentoring commitment · Train-the-trainer · Active SRE peer network

◆ WHY "PROPOSED" IS ESSENTIAL — THE CORPUS IS NOT A DECREE

The nomenclature only exists if organisations choose to adopt it. This choice is pragmatic, not ideological. An organisation that adopts it gives itself a more precise competence management tool than "DevOps" — nothing more. The gentle exit begins with an internal decision, without waiting for any collective agreement or branch decree.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

WHAT EXISTS TODAY — WITHOUT WAITING FOR ANY NEW INSTITUTION

Whatever the decision on the nomenclature, the training tools exist today. They require no collective agreement, no branch accord, no new institution. They require an HR decision and a training budget — available in any organisation that uses CPF or equivalent professional development funding.

◆ NEUTRAL CERTIFICATIONS — AVAILABLE, PORTABLE, LESS EXPENSIVE

Linux Foundation: LFCA (entry, €395), LFCS (systems administrator, €395), LFCE (engineer), CKA (Kubernetes administrator, €395), CKAD, CKS (security). Practical hands-on exam — no multiple choice. Real competence or the exam fails. CPF-eligible. Less expensive than equivalent AWS certifications (€300-800).

HashiCorp: Terraform Associate, Vault Associate. Neutral tools, multi-cloud, essential for portable IaC — no hyperscaler dependency.

CNCF: cloud-native neutral certifications, based on open source standards. No vendor lock-in.

These certifications produce transferable competences across environments — exactly what "The CPF Misused" calls for funding instead of hyperscaler cloud certifications that reinforce lock-in.

◆ THE FOUNDATION OF IRON — A PROPOSED STRUCTURED REBUILDING PROGRAMME

The Foundation of Iron is a corpus proposal (26 weeks, electricity to Kubernetes) — not an existing institutional programme. It is a pedagogical plan that can be implemented by any trainer or training organisation. For the converting DevOps, the path is targeted, not linear:

Modules 1-6 (Electricity, Hardware, Physical Server): mandatory physical layer for all levels. Missing in 90% of DevOps trained solely in cloud.

Modules 7-12 (Linux fundamentals, Storage, Virtualisation): deepening what the DevOps often has only at surface level.

Modules 13-20 (Physical OSI networking, Routing, Active Directory): the real differentiator between a DevOps and a SysOps/NetOps L4.

Modules 21-26 (Network security, Architecture, Bare-metal Kubernetes): L5-L6 level deepening.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · SUPERVISED ROTATION AND MENTORING — AVAILABLE TODAY

SIX MONTHS ALONGSIDE A PRINCIPAL SRE — WHAT NO COURSE REPLACES

Theoretical training gives structure. Contact with the machine and with the person who masters it gives flesh. This mechanism needs no new institution — it exists in every organisation that still has an experienced Principal SRE. It only needs to be formalised: a written mentoring agreement, defined objectives, a time limit.

◆ THE OBSERVATION DAY — THE FIRST STEP AT ZERO COST

Documented in "The Mentor and the Metal": one day without assessment, without badge, without PowerPoint. The converting engineer observes. They accompany the Principal SRE during maintenance interventions, a component replacement, an incident if the occasion arises. They eat with the team. They physically see what their architecture diagrams describe from the outside.

Cost: one day of the Principal SRE's time in accompaniment mode. Available tomorrow, in any organisation, with no exceptional budget.

◆ STRUCTURE OF THE FORMALISED 6-MONTH ROTATION

Months 1-2 — Immersive observation: accompaniment without intervention. Documentation, questions, reading runbooks in real situations. The Principal SRE explains decisions in real time.

Months 3-4 — Supervised practice: interventions under direct supervision. Cabling, configuration, diagnosis. The "why" behind each action — tacit knowledge that cannot be written but can be transmitted.

Months 5-6 — Supervised autonomy: autonomous interventions on non-critical systems. The Principal SRE is available but not present.

Minimal formalisation: a one-page mentoring agreement — objectives, duration, evaluation method. Referenced in both parties' job descriptions. Not a new institution: a professional commitment between two people, recognised in their respective responsibilities.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE NOMENCLATURE AS AN ENTRY DOOR FOR WOMEN

A BODY WITHOUT A NAME HAS NO VISIBLE DOOR — WHICH EXCLUDES WOMEN

There is a connection nobody has explicitly documented yet: a clear nomenclature is a necessary condition for women to enter the infrastructure body. Not sufficient — but necessary. "The Invisible Amputation" documented that women represent 10% of the infrastructure body where 50% are needed. This study asks the mechanism question: why so few? And how the proposed nomenclature is part of the answer.

◆ AMBIGUITY AS AN INVOLUNTARY EXCLUSION MECHANISM

A woman looking at the infrastructure job market today sees: "DevOps", "SRE", "Cloud Engineer", "Platform Engineer", "DevSecOps", "Site Reliability", "Infrastructure Engineer" — vague titles, no clear trajectory, no readable progression grid, no objective signal on what is expected at each level.

This ambiguity discourages profiles who do not have a network to decode it. Women, statistically less present in the informal networks of infrastructure (tech evenings, conferences, non-mixed Slack channels), bear this decoding burden more than men who can access it through peers. The ambiguity is not intentionally sexist — but it produces real exclusion.

◆ THE PROPOSED NOMENCLATURE AS A READABLE DOOR

A clear nomenclature says to the woman looking from the outside: here is the path, here are the objective criteria at each level, here is where one can go. No decoding through informal networks. No "you need to know someone who knows someone". The criteria are public, the levels are readable, the progression is documented.

The proposed nomenclature — if adopted by even a few organisations — would immediately produce two effects: job offers with precise titles and objective criteria (attracting profiles without networks), and visible progression grids (enabling projection without having to ask a peer). These are the two minimum conditions for a woman to envisage a path in infrastructure without having to guess the implicit rules.

◆ THE CAUSAL CHAIN — NOMENCLATURE → DOOR → WOMEN → COMPLETE BODY

A named body creates a visible door. A visible door creates a possible entry. A possible entry triggers the arrival of those who were waiting for a readable signal. The body becomes complete. This is the reverse mechanism of the Invisible Amputation — and it begins with an act as simple as adopting a nomenclature in a job description.

HUMAIN

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE HONOURABLE EXIT FOR ORGANISATIONS

REBUILDING WITHOUT RAZING — OR PUBLIC ADMISSION OF ERROR

The organisation that built its teams on the DevOps wave did not make an error — it made a coherent choice with the information available at the time. The honourable exit does not ask for a public admission of fault. It proposes a progressive, documented transition plan using the tools that exist today.

◆ THE TRANSITION PLAN IN FOUR PHASES

Phase 1 — Internal diagnosis (months 1-2): map real competences of each team member using the proposed nomenclature as a reading grid. Tool: structured interview between the Principal SRE and each engineer. Result: individual gap map and targeted training plan. Cost: the Principal SRE's time. Budget: zero.

Phase 2 — Targeted training (months 3-9): Foundation of Iron targeted modules + Linux Foundation certifications according to target level. Budget: €1,200 to €2,400 per person. CPF can fund Linux Foundation certifications — they are eligible. Less expensive than AWS. No exceptional budget needed.

Phase 3 — Progressive renaming (months 6-12): if the organisation decides to adopt the proposed nomenclature, update job descriptions. Without redundancy. With salary valorisation. Renaming follows training — not the reverse. Each title changed is a door opened toward profiles who could not project themselves into "DevOps".

Phase 4 — Consolidation (months 12-18): Technical Primacy clause in the Principal SRE's contract, formalised mentoring in senior job descriptions, job offers with proposed nomenclature titles. The nomenclature becomes an internal reality — without waiting for it to be an external norm.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

TWO SIGNALS THAT CREATE MOVEMENT — THE EXAMPLE AND THE LEGIBILITY

The corpus does not need everyone to be convinced. It needs two things simultaneously: an ambassador who made the journey and tells it, and a readable door for those who want to enter. The proposed nomenclature creates the door. The ambassador creates the desire to walk through it.

◆ THE AMBASSADOR — THE ONE WHO MADE THE JOURNEY

The DevOps who completed six months of supervised rotation, passed the LFCS, changed their title to Lead SysOps in their job description, is the best ambassador for the proposal. "The Uniform of the Body" documented the chain of gaze: the one who changed level gives others a concrete reason to make the journey. The ambassador does not convince through theoretical argument — through visible example and proof that the path exists and leads somewhere.

The concrete testimony: "I was a DevOps for eight years. Six months of supervised rotation with our Principal SRE showed me I did not understand what I was deploying. I passed the LFCS. My title changed. My salary was revised. And my legitimacy in technical arbitrations is now recognised." This testimony is worth ten studies. The corpus creates the conditions. The ambassador creates the movement.

◆ THE DOOR — WHAT THE NOMENCLATURE CREATES FOR WOMEN

When an organisation publishes a job offer "SysOps L3 — criteria: LFCS or equivalent, 2 years Linux systems experience, knowledge of OSI layers 1-4" — this offer is readable by a woman who has no network to decode "Senior DevOps with 3-5 years XP and full-cloud stacktrace". Objective criteria replace implicit decoding. The door is open — not just drawn on a wall.

This is the reverse mechanism of the Invisible Amputation. Not a diversity declaration of intent. Not a quota. A clear nomenclature, objective criteria, readable offers. The body becomes complete because the door exists — not because someone decided to force women through it.

HUMAN

The gentle exit does not say: you were wrong.
It says: here is what is missing — with today's tools.
It says: here is the door — it is open for everyone.
Come.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE ACCOMPANIED EXODUS

Managed Services as the Exit Door from Hyperscaler Cloud

Migration · Transmission · Progressive Sovereignty

◆ THE MISSING PIECE OF THE CORPUS

The Opération Dindon corpus documented the problem (lock-in, loss of bare-metal), the causes (cognitive capture, infrational loop), and the tools (Sovereign Interface, Gentle Exit). What was missing was the operational entry point for organisations that want to act now but no longer have anyone who knows how to manage bare-metal. Managed services are that door. They do not give sovereignty immediately — they create the conditions for it to become possible, accompanying the organisation throughout the exodus journey.

Note: the operational elements of this study draw on Amine RAITI's lived experience during his time at Ecritel, cited here as a concrete example and not as advertising. The approach may differ between managed service providers.

SPECTRUM	CONCEPT	WATERMARK
HOUSING → FULL MSP	**ACCOMPANIED EXODUS**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · THE PROBLEM MANAGED SERVICES SOLVE

THE ORGANISATION WANTS TO LEAVE THE CLOUD · IT NO LONGER HAS ANYONE WHO KNOWS BARE-METAL

The corpus documented the "Gentle Exit" as an architectural principle. It documented the "Sovereign Interface" as a technical tool. These two studies answer "how" to exit. They do not answer "with whom" for organisations that no longer have the internal skills to do it.

◆ THE TYPICAL PROFILE OF THE BLOCKED ORGANISATION

It migrated to cloud 5-8 years ago. It let go of or lost its system engineers, network administrators, DBAs. It has a competent DevOps team on AWS or GCP — but with no bare-metal culture. It pays growing egress fees. It read the Terms of Service and discovered California jurisdiction. It wants to leave.

It cannot. Not alone. Its DevOps know how to deploy on EKS. They do not know how to configure a physical switch, dimension an IP addressing plan, choose a storage array, install a hypervisor, configure bare-metal monitoring. The competence left with the engineers who departed. This is "The Departure of the Last One Who Knows" applied at organisational scale.

◆ WHAT MANAGED SERVICES BRING THAT CORPUS STUDIES DID NOT

The Sovereign Interface is an architecture. It does not implement itself — engineers are needed to write it.

The Gentle Exit is a principle. It does not execute itself — engineers are needed to migrate.

The Foundation of Iron is training. It takes 6 to 26 weeks — and the organisation does not have 26 weeks of runway before its next AWS bill.

The managed service provider is the missing engineer. They have the bare-metal competence the organisation abandoned. They know the migration path because they have done it dozens of times. And — fundamental difference from a hyperscaler — their business model is not to create a new dependency but to make the organisation capable.

◆ NASSIHA — MANAGED SERVICES ARE NOT A NEW LOCK-IN

The difference between a sovereign managed service provider and a hyperscaler: a managed service contract can be terminated with reasonable notice (generally 3-6 months). There are no "noncancellable" commits. No California jurisdiction. No egress fees. And if the contract is well drafted, the documentation and knowledge transmitted remain with the client at the end.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

HOUSING · PARTIAL MANAGED SERVICES · FULL MANAGED SERVICES

LEVEL I — HOUSING (COLOCATION)

What the client brings: their own racks, their own servers, their own complete technical skills.

What the provider brings: the datacentre, redundant power, cooling, fibre optic, physical security, 24/7 building monitoring.

Who manages what: the client manages everything from their own teams. The provider does not touch client machines unless explicitly requested.

Profile: organisations that still have bare-metal engineers but no own datacentre. The most sovereign form — the client masters everything except the building.

Risk: if bare-metal engineers leave, the client is left with hardware they can no longer manage.

LEVEL 2 — PARTIAL MANAGED SERVICES

What the client retains: control of the application layer, databases, logical network, deployments. The top of the stack.

What the provider manages: hardware, OS, virtualisation, physical network, low-level monitoring, backups, security updates.

Profile: organisations with DevOps or developers, but without system administrators or physical network engineers. The most common model after a cloud migration.

Added value: the client progressively builds competence on lower layers, with the provider's engineers available to explain and train.

LEVEL 3 — FULL MANAGED SERVICES

What the client retains: defining needs, validating SLAs, monitoring dashboards, business decisions. Vision — not execution.

What the provider manages: everything, A to Z. Hardware, network, OS, middleware, monitoring, backups, security, updates, daily operations.

Profile: organisations with only business profiles or cloud-native DevOps with no operations culture. The Accompanied Exodus starting point.

Objective: full managed services is not a destination — it is a transition phase. The goal is to progressively move toward partial managed services or housing as teams build competence.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

WHAT ACTUALLY HAPPENS WHEN A MANAGED SERVICE PROVIDER WINS A CLIENT

What follows is based on lived experience at Ecritel. This is not advertising — it is a concrete example. The approach may differ between providers, but the main steps are structurally similar.

◆ STEP 1 — THE THREE-PARTY KICKOFF MEETING

Upon contract signature, a meeting brings together three parties: Ecritel teams (migration project manager + system engineers), the previous host's teams, and the client's teams. The objective is not to criticise the existing setup — it is to understand it precisely.

Two scenarios for access to the existing infrastructure:

Scenario A — Export by the previous host: the previous host produces data exports (databases, file systems, configurations), network diagrams, service lists. Ecritel receives and works from the exports.

Scenario B — Direct root access: when the previous host does not cooperate or when the timeline requires it, Ecritel teams request direct root access to the machines and carry out the inventory and exports themselves. They take operational responsibility for operations performed on the existing infrastructure.

◆ STEP 2 — INVENTORY AND ANALYSIS OF THE EXISTING SETUP

The inventory covers all OS types (Linux distributions, Windows Server, BSD as applicable), all databases (MySQL, PostgreSQL, Oracle, SQL Server, MongoDB, Redis — each engine has its own export and migration procedures), middlewares, network configurations (VLANs, firewall rules, load balancers), certificates, cron jobs, operational scripts. Nothing is assumed — everything is verified and documented.

The analysis of the existing setup produces a dependency map — which services depend on which others, in what order they must be migrated to avoid breaking the chain.

HUMAN

◆ STEP 3 — PHASED MIGRATION PLAN AND FOLLOW-UP MEETINGS

Migration does not happen overnight. It is phased according to dependencies and risks. Each phase has a schedule, success criteria and a rollback procedure. Follow-up meetings are organised throughout the project — weekly during active migration, bi-weekly during stabilisation. The client monitors progress. They test proper functioning step by step. They validate before moving to the next phase. This validation is not passive — it is on-the-job training.

◆ STEP 4 — TAD AND INFRASTRUCTURE DOCUMENTATION

At the end of the migration, Ecritel teams write the Technical Architecture Document (TAD) — an exhaustive photograph of the infrastructure put in place: network diagrams, detailed configurations, operational procedures, backup and restore procedures, escalation contacts. This document belongs to the client. It is the first Runbook of the new infrastructure. And depending on the contract type, Steering Committees (COPILs) are organised — at Ecritel or at the client — to present SLAs, performance indicators, monthly incidents, planned developments.

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · SKILLS TRANSMISSION — COMPAGNONNAGE APPLIED

THE PROVIDER IS NOT THERE FOREVER — THEY ARE THERE SO YOU NO LONGER NEED THEM

The fundamental difference between a sovereign managed service provider and a hyperscaler is not technical — it is contractual and intentional. A hyperscaler wants you dependent on them forever. A good managed service provider wants you to build enough competence that you no longer need them — or need them only for high-value-added operations. This is not philanthropy — it is a different business model, based on trust and relationship longevity rather than lock-in.

◆ TRANSMISSION IN PRACTICE — WHAT ACTUALLY HAPPENS

Documentation as a transmission act: the TAD written by the managed service provider is not a bureaucratic document — it is the first act of transmission. The client's DevOps who read the TAD learn how their infrastructure works. Those who participate in writing it learn even faster.

COPILs as disguised training: when the provider presents monthly SLAs, they explain what happened, why, how it was resolved. The client's technical teams who attend these meetings are in continuous training without it being called training.

Shared on-call: in some contracts, the provider's engineers are available to the client's teams during incidents — not to replace them but to accompany them in resolution. A 3am incident resolved with the provider's engineer online teaches more than a full day of classroom training.

Formal knowledge transfer: in well-drafted contracts, a knowledge transfer phase is explicitly planned — training sessions on the client's specific infrastructure, handover of operational procedures, support for the first operations performed autonomously by internal teams.

◆ DIGITAL COMPAGNONNAGE APPLIED AT ORGANISATIONAL SCALE

The "Digital Craftsmen" model proposed in the corpus describes an aspirant learning alongside a maâllem. Partial managed services with knowledge transmission is this model applied at the scale of an entire organisation. The provider is the collective maâllem. The client's DevOps are the aspirants. The migration and daily operations are the joint masterpiece. And at the end — if everything went well — the client has a team that knows how to manage its infrastructure and a provider who can focus on higher-value-added operations.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE PROVIDER ADAPTS TO YOU · NOT THE OTHER WAY AROUND

◆ OPERATIONAL FLEXIBILITY — THE HAPROXY/FTP EXAMPLE

Lived experience at Ecritel: a client has two machines — prod and preprod. They want to promote preprod to prod without transferring code, simply by swapping roles. Solution proposed: a file deposited via FTP by the client containing `machine1=prod ; machine2=preprod`. A script reads the file every X minutes and automatically reconfigures the HAProxy backend. The client remains in full managed services but is completely autonomous on deployments, with no provider intervention required per cycle.

What this illustrates: the provider wrote the script. They adapted the infrastructure to the business need. The opposite of cloud: "adapt your architecture to our managed services."

◆ REAL INTERNATIONAL PRESENCE · INCLUDED TRAFFIC · DEDICATED FIBRE

Ecritel accompanied clients to obtain a CSP licence in China from Paris. Infrastructures were synchronised between China and France via Unison or lsync for images. Pseudo-L2 configurations at international scale were set up to ensure network continuity between distant sites. Dedicated fibre links can connect client offices or another site — something cloud cannot do physically. Traffic is included in the contract or at very reasonable rates. No egress fees discovered at end of month.

◆ SUPPORT THAT TALKS TO A TECHNICIAN — NOT A TICKET

You pick up the phone. A technician answers. They know your infrastructure. They can act immediately or escalate to an engineer — included in the standard contract, not in an overpriced Premium plan. For housing clients, physical proximity actions are possible: a technician can go into the server room to cable, reboot, inspect. This is not comparable to AWS support — form, ticket, delay measured in days, Premium plan at tens of thousands of euros per year to get a human on the phone.

HUMAN

◆ THE BCP — THE MOST POWERFUL ARGUMENT FOR SOVEREIGN MANAGED SERVICES

The Business Continuity Plan is the argument that CIOs understand immediately — because it is their personal responsibility when infrastructure goes down.

The hyperscaler BCP is a software architecture. Multi-AZ, multi-region, automatic failover — sold as a turnkey solution. In reality: dozens of proprietary services to master, a cost of 2x to 5x the primary infrastructure, and a failover test that nobody really does because it is expensive and frightening. When the real incident arrives at 3am, DevOps discover that the failover site's Terraform has not been updated in 8 months.

The managed service provider BCP is a muscular reflex. The provider has two datacentres connected by a dedicated fibre they physically control. They have performed hundreds of real failovers — not sandbox tests. Their team knows the friction points, the services that always take 3 extra minutes to restart. They have a Runbook for every incident type. And they adapt the BCP to the client's real application needs: RPO, RTO, synchronous or asynchronous replication per service. Not a checkbox in a catalogue — a bespoke architecture.

The formula: the hyperscaler BCP is an architecture. The managed service provider BCP is a muscular reflex. One is documented in a Confluence nobody reads. The other is repeated until the team can do it with their eyes closed. The corpus study "The Gentle BCP" documents this mechanism in detail.

◆ THE HUMAN RELATIONSHIP AS A DURABLE ASSET

Cloud gives you an SLA. A managed service provider gives you a relationship. Over time, the relationship is worth more than the SLA.

The technician who answers at 3am knows your infrastructure, your business constraints, sometimes your DBA's first name. The quarterly COPIL is not a contractual meeting — it is a conversation between people who know each other. Over five years, this mutual knowledge produces what no SLA can produce: trust.

A client who trusts their provider alerts them to an anomaly before it becomes an incident. They call before making a risky architectural decision. On the other side, the engineer who knows their client does a little more than the contract requires — because the relationship has a value they want to preserve.

Some of my clients at Ecritel are still my friends today. This bond was not created by the contract — it was created by contact. The COPILs, the 3am calls, the visits to the client. Managed services done well does not produce contracts. It produces relationships. And relationships outlast contracts.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · SOVEREIGN MANAGED SERVICE PROVIDERS — WHO CAN DO IT

ECRITEL · DRI · OVH SUNRISE · SCALEWAY · CLARANET · INFOMANIAK — THEY EXIST AND THEY KNOW HOW

◆ WHAT DISTINGUISHES A SOVEREIGN MSP FROM A HYPERSCALER MANAGED SERVICE

Managed AWS / Google Managed Services / Azure Managed Services: they manage your cloud infrastructure for you — but on AWS, GCP or Azure. Dependency on the hyperscaler is maintained, sometimes reinforced. California jurisdiction remains. Egress fees remain. Application lock-in remains. These are not sovereign MSPs — they are subcontractors of the hyperscaler.

Sovereign MSPs (Ecritel, DRI, OVH, Scaleway, Claranet, Infomaniak): they manage your infrastructure on their own datacentre or neutral infrastructure, under French or European law. Their contracts are terminable. Documentation produced belongs to you. Data stays in Europe under European jurisdiction. And their certifications (HDS for health, SecNumCloud for critical infrastructure operators, ISO 27001 for general security) cover regulatory requirements that hyperscalers struggle to fully satisfy.

◆ THE CERTIFICATIONS THAT MATTER

HDS (Health Data Hosting): mandatory for hosting health data in France. Major French MSPs are HDS-certified. AWS and GCP are too — but under contractual conditions that may conflict with GDPR and the CLOUD Act.

SecNumCloud (ANSSI): the French cloud sovereignty label. Requires that operating companies are subject to French law, with no possibility of extra-European takeover. AWS and GCP cannot obtain it. This is the label that makes the difference for critical infrastructure operators and sensitive public administrations.

ISO 27001: international information security management standard. Most serious MSPs are certified. Minimum condition for public tenders and major accounts.

◆ SELECTION CRITERIA FOR AN MSP FOR THE EXODUS

Three questions to qualify a managed service provider:

1. **Do you have documented experience of migrations from AWS/GCP/Azure?** (Ask for references.)
2. **What is your documentation policy?** (Does the TAD and all documentation produced belong to the client from the end of the engagement?)
3. **How do you organise knowledge transfer to our internal teams?** (If there is no precise answer, this is a warning signal.)

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

WHAT THE CLIENT MASTERS AT THE END OF EACH PHASE

PHASE 1 — INVENTORY & MAPPING (1-2 MONTHS)

Three-party kickoff. Root access or exports from previous host. Exhaustive inventory: OS, databases, middlewares, network, certificates, cron jobs, scripts. Dependency mapping. Migration risk identification.
What the client masters at the end: a faithful and complete picture of their current infrastructure — often for the first time.

PHASE 2 — TARGET ARCHITECTURE & MIGRATION PLAN (1 MONTH)

Definition of target architecture on new infrastructure. Phased migration plan with migration order by dependencies. Rollback procedures for each phase. Client validation before start.
What the client masters at the end: what their infrastructure will look like — and why each choice was made.

PHASE 3 — PHASED MIGRATION (3-12 MONTHS DEPENDING ON COMPLEXITY)

Service-by-service migration per the validated plan. Client tests proper functioning at each step before validation. Weekly follow-up meetings. Managing migration incidents (there will be some). Plan adjustments as needed.
What the client masters at the end: their infrastructure runs on the new platform. DevOps have followed each migration and begin to understand the lower layers.

PHASE 4 — TRANSMISSION & SKILLS BUILD-UP (6-12 MONTHS IN PARALLEL)

TAD finalised and handed to client. Monthly or quarterly COPILs per contract. Formal knowledge transfer: training sessions on the client's specific infrastructure, support for first autonomous operations. Progressive transition toward partial managed services then housing as teams develop.
What the client masters at the end: a documented infrastructure, a team building competence, a decreasing dependency on the provider. Progressive sovereignty is underway.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

8

SECTION 8 · MANAGED SERVICES IN THE CORPUS — THE MISSING PIECE

EVERY CORPUS STUDY WAS WAITING FOR THIS ANSWER WITHOUT NAMING IT

CORPUS STUDY	WHAT MANAGED SERVICES BRING IN RESPONSE
The Gentle Exit	Architectural principle → Concrete operation. The MSP executes the gentle exit.
Departure of the Last One Who Knows	The MSP is temporarily "the last one who knows" — but their contract requires them to transmit before leaving.
The Digital Craftsmen	Compagnonnage at organisational scale. MSP = collective maâllem. DevOps = aspirants.
The Foundation of Iron	26-week classroom training → on-the-job training in production during migration. More effective.
The Sovereign Interface	Architecture → Implementation. The MSP can deploy it during migration.
No Sovereignty Without Matter	Aspiration → Path. Managed services is the most accessible route to material sovereignty.
أتستبدلون	The answer to the question. You do not trade what is better for what is inferior — you return to what is better via the MSP.
The Replacement that Reveals	The MSP reveals what the organisation no longer knows how to do — and transmits what is needed to know it again.

HUMAN

The Accompanied Exodus is not a flight — it is a return.
Return to mastery of the substrate. Return to sovereignty.
Return to the ability to say no.
The managed service provider does not make the journey for you — they make it with you,
until you find your legs to make it alone.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

MAJOR STUDY · OPERATION DINDON · JUNE 2026 · LIBERATION ARCHITECTURE

THE SOVEREIGN EXODUS

Global Architecture for Digital Lock-in Exit · From the Maths Teacher to the Boardroom

6 Strata · 3 Horizons · School · Training · Organisation · Hardware · Legal · Institutional

HUMAN

◆ THE THESIS — FROM AUDIT TO ACTION

The Operation Dindon corpus documented captivity. This study documents the exit. Not a wish list — an operational liberation architecture across six simultaneous strata and three time horizons. Every lever is immediate, fundable, and technically achievable. The bifurcation at school costs almost nothing. The bifurcation at t=20 years costs a total rewrite. This study begins with school.

PAGES	STRATES	HORIZONS	WATERMARK	CORPUS REF.
16	6	3	HUMAN DOMINANT	64

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

CC BY-NC-SA 4.0 · *Operation Dindon · June 2026 · Ultimatum: 8 September 2026*

LOCK-IN EXIT IS NOT A PROJECT — IT IS A SIMULTANEOUS ARCHITECTURE ACROSS SIX FRONTS

Digital perdition is a seven-node causal chain documented in Anatomy of Digital Perdition (16p). The exit is not breaking a single link — it is the simultaneous reconstruction of six strata. None is sufficient alone. Together, they create an irreversible bifurcation.

STRATUM	⚡ IMMEDIATE	🔧 MEDIUM TERM	🏛️ LONG TERM
1 · School	Maths teacher → dish + RPi	Physics · Tech · Art	Official Anti-Amputation Core
2 · Training	56h Linux/Network · neutral CPF	Women 30+ · RNCP · Job codes	National engineering school plan
3 · Organisation	Arch. question · Job titles	Portability Score JIRA · Cap. Planning	Cloud-free BCP · Peer mentoring
4 · Hardware	Grade A Refurbished + TPM	Accompanied Exodus · portable workloads	RISC-V · ARM · Sovereign fabs
5 · Legal	Ultimatum · Terms/CLOUD Act training	Sovereign specs · Reversibility clause	Egress fee law · Insurable SecNumCloud
6 · Institutional	Balanced CPF · ESN partners	Sovereign DINUM · Public Cloud	Chips Act · FFA insurance · ANSSI

◆ GUIDING PRINCIPLE — EARLY BIFURCATION

Each stratum can be activated independently. The effect is cumulative — an organisation activating strata 3 and 4 without waiting for strata 1 and 2 still progresses. Complete exit is a simultaneous architecture. Partial exit is preferable to paralysis. cf. Anatomy of Digital Perdition — Section 6: The Law of Late Bifurcation.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

I

STRATUM I · SCHOOL — THE MATHS TEACHER AS HEAD OF SRE

CUTTING THE CAUSAL CHAIN AT THE ROOT — BEFORE FREE TIER CORRUPTS THE STUDENT'S SYSTEM 1

◆ ⚡ IMMEDIATE — THE SATELLITE DISH, THE RASPBERRY PI, THE CNC

The mathematics teacher teaches sine and cosine. Before drawing the curve on the board, they bring out a satellite dish and say: 'Look how a wave propagates in physical reality.' No ministerial budget required. A satellite dish, a €35 Raspberry Pi, and the awareness that physical substrate exists.

This is the school's Head of SRE — not because they manage infrastructure, but because they are first to show that mathematical abstraction has a physical body. Sine and cosine exist in a satellite antenna. A network exists in cables. A server is controlled electricity. This awareness at age 14 is the vaccine against Free Tier.

◆ 🔧 MEDIUM TERM — THE COMPLETE PEDAGOGICAL CHAIN

The physics teacher takes over — showing the electromagnetic laws driving the CNC motor. From Faraday's law to the rotor: G-code moves electrons, electrons move the magnetic field, the field moves the motor, the motor moves the tool, the tool machines matter. The complete causal chain from symbol to matter.

The technology teacher shows the CNC — how a text file becomes a metal part at 0.01mm precision. Same logic as the LLM: input → processing → output. The Mystical Machine is demystified before marketing sacralises it.

The art teacher shows the 3D printer — how an STL file becomes a tangible object layer by layer. Creativity has a physical substrate. The digital touches matter.

◆ 🏛️ LONG TERM — THE OFFICIAL ANTI-AMPUTATION CORE · CF. THE MATHS TEACHER'S BAG (15P)

An official programme in schools: 4 hours per term of 'digital physics' — CNC · 3D printer · physical network · server. Estimated budget: €500 per school for hardware (Raspberry Pi, cables, switch). Not an IT budget — a pedagogical budget. Teacher training: 2 days. National deployment: 3 to 5 years.

HUMAN

THE COMPLETE PEDAGOGICAL CHAIN — FROM SYMBOL TO MATTER · EVERY TEACHER IS A LINK

◆ WHY FOUR DISCIPLINES AND NOT ONE

The scholastic amputation documented in The Maths Teacher's Bag (15p) does not come from a single course — it comes from the absence of coherence between disciplines. The mathematician teaches abstraction. The physicist teaches law. The technician teaches gesture. The art teacher teaches form. None shows the chain connecting them. This continuity rupture creates the cognitive void that Free Tier fills.

◆ THE INTEGRATED PEDAGOGICAL PROJECT — ONE 3D PRINTER AS THREAD

Concrete and immediately deployable proposal: an annual pedagogical project around a 3D printer.

Maths: geometric modelling → STL file. Physics: extrusion thermodynamics, motor electromagnetism. Technology: Marlin firmware, slicing parameters, G-code file. Visual arts: object design, formal constraints. IT: printer network connection, Octoprint server on Raspberry Pi.

In one project, the student understands the digital has a body, code has an energy cost, abstraction has a physical substrate. And that nobody needs AWS to print a plastic object.

◆ THE REAL COST — WHAT IS POSSIBLE NOW WITHOUT OFFICIAL PROGRAMME

A budget FDM 3D printer: €250. A Raspberry Pi 4: €55. An 8-port network switch: €25. Cables: €15. Total: €345 per school. A parent-teacher association can fund this. A local authority can fund this. A local company donation can fund this. This project requires no ministerial decision — a motivated teacher and €345. The cheapest bifurcation in the corpus.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

TURNING THE WEAPONS OF CAPTURE AGAINST THEMSELVES — CPF CAN FUND FREEDOM

◆ ⚡ IMMEDIATE — THE 56 HOURS AS DUNNING-KRUGER BREAKER

56 training hours per year in certain conventional frameworks. Sufficient volume to acquire the fundamentals that break the cloud mastery illusion from IaC:

Week 1 (14h): Fundamental networking — OSI model · ARP table · VLAN · IP routing · TCP/IP. What Lambda never shows.

Week 2 (14h): Fundamental Linux — filesystem · inodes · processes · service management · systemd. What Kubernetes hides.

Week 3 (14h): Storage — RAID · LVM · snapshots · physical backup. What S3 replaces without anyone understanding what it replaces.

Week 4 (14h): Physical security — certificates · PKI · firewall · ACL. What 'managed IAM' abstracts.

In 56 hours, no senior is produced. Dunning-Kruger is broken. The developer who has seen an ARP table will never again confuse an HCL script with infrastructure mastery.

◆ 🔧 MEDIUM TERM — THE REVERSED CPF · NEUTRAL CERTIFICATIONS IN CATALOGUE

CPF massively funds hyperscaler certifications (documented in The CPF Misused, 6p). The answer is not elimination — it is balance. Massively reference neutral certifications in the CPF catalogue:

LFCS (Linux Foundation Certified System Administrator) · RHCSA (Red Hat Certified System Administrator) · CompTIA Network+ · CompTIA Security+ · LPIC-1/2 · Kubernetes CKA (neutral, not AWS proprietary)

These certifications exist. They are not visible enough in CPF catalogues. A France Compétences decision — not a technical one. RNCP registration timeline: 6 to 18 months. Main actors: France Compétences + Ministry of Labour.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

RETRAINING AFTER 30 — THE MOST UNDERESTIMATED LEVER IN THE CORPUS

◆ WOMEN RETRAINING AFTER 30 — WITHOUT PROPRIETARY REFLEXES

A woman retraining in IT at 32 does not have the proprietary reflexes of a developer trained on AWS Educate at 20. Nobody offered her Free Tier at school. No AWS certification to defend. No ego invested in a 5-year cloud architecture.

She learns on fundamentals — because that is what retraining programmes teach when well designed. The documentable trajectory: technical support levels 1 and 2 → systems administrator → infrastructure engineer → senior SRE. This trajectory does not go through AWS Educate. It goes through Linux, networking, and systems physics.

Training cost: fundable on CPF + Pôle Emploi/France Travail retraining aids. Time to first position: 6 to 12 months for support. Time to full autonomy: 3 to 5 years. The same timeline as any technical track.

◆ THE HR NUDGE — RNCP AND ROME JOB CODES AGAINST ATS SABOTAGE

The HR lock documented in Anatomy of Digital Perdition S1B (−34% Linux positions in 4 years): ATS filters on proprietary cloud keywords. The answer has two levels:

Level 1 — RNCP: register LFCS, RHCSA, CompTIA Network+ in the RNCP (National Professional Certification Registry). An RNCP certification is legally enforceable against an ATS. France Compétences is the actor. Timeline: 6-18 months.

Level 2 — ROME job codes: modify Pôle Emploi's ROME (Operational Repository of Occupations and Jobs) to reintegrate 'Infrastructure Engineer', 'Systems Administrator' and 'Network Engineer' as distinct occupations from 'Cloud Engineer' or 'DevOps Engineer'. A ROME code is the reference HR uses to build job descriptions. Actors: DARES + Pôle Emploi.

◆ PEER-TO-PEER MENTORING — DISMISSED SENIORS AS TRANSMITTERS

Senior infrastructure engineers dismissed between 2018 and 2023 under the NoOps pretext still exist — freelancing, in early retirement, in niche ESNs. They carry 20 to 30 years of low-level skills that nobody transmits. A peer-to-peer mentoring programme — seniors → retrainees, especially women after 30 — is the fastest and least costly lever. Platform: France Travail + professional associations. Budget: near zero. The corpus is the reference. Seniors are the transmitters.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

REBUILDING THE INFRA/DEV BOUNDARY — THE SOVEREIGN ORGANISATIONAL API

◆ ⚡ IMMEDIATE — THE ARCHITECTURAL QUESTION AND TITLE REHABILITATION

The architectural question: 'Can this component run outside AWS without a rewrite?' asked before every infrastructure technical decision. Cost: zero. Timeline: immediate. It makes conscious the decision that was unconscious (Default Effect, Stratum 6 of The Architecture of Consent).

Title rehabilitation: 'Infrastructure Engineer', 'Systems Administrator' and 'Network Engineer' must reappear in org charts and budgets as distinct titles — not subsets of a generic DevOps role. What has no budget line does not exist in the boardroom. What has no title is not recruited. Actors: HR + CIO. Timeline: next budget cycle.

◆ 🔧 MEDIUM TERM — THE PORTABILITY SCORE IN JIRA AND THE DEFINITION OF DONE

The Product Owner is evaluated on velocity. Architectural reversibility is perceived as a sprint brake. The answer is not to block the PO — it is to make portability debt as visible as code technical debt.

The Portability Score (0 to 3):

0 = Pure proprietary (native DynamoDB, Lambda with AWS LAYERS, Cognito) — no migration possible without total rewrite

1 = Migration possible with major effort (12-36 months)

2 = Migration possible with moderate effort (3-12 months)

3 = Agnostic — runs on any infrastructure without rewrite

This score is a JIRA ticket field. Visible in the Definition of Done. Does not block the sprint — makes debt visible to management, CFO, and infra engineer. Setup cost: JIRA configuration, 2 hours. Deployment: immediate.

◆ THE SOVEREIGN IDP — THE INFRA ENGINEER AS BRICK PROVIDER · cf. The Dissolution of the Boundary (8p)

The sovereign Internal Developer Platform (IDP): the infra engineer provides agnostic, reversible bricks — a database brick (PostgreSQL bare-metal or cloud), a message queue brick (RabbitMQ or SQS depending on context), an authentication brick (Keycloak or Cognito). Each brick must answer 'yes' to the architectural question. Training to design a sovereign IDP: 3 to 5 days on OpenStack, bare-metal Kubernetes, PostgreSQL, standardised APIs. Fundable on CPF.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

REBUILDING THE FUNCTIONS THAT NOOPS DESTROYED — WITHOUT RECREATING THE DEVS/OPS WAR

◆ CAPACITY PLANNING AS MANDATORY FUNCTION

The disappearance of Capacity Planning (documented in NoOps · Autopsy, 6p) produced uncontrolled cloud invoices. Restoring it does not immediately require hiring a dedicated expert — it requires a rule: no production deployment without a real-scale monthly cost estimate signed by an infrastructure engineer. Not bureaucracy — translating software architecture into real costs before invoicing. The formula is documented in The Mystical Machine S5 (TCO + PUE). Applicable within a week.

◆ QUARTERLY IaC AUDIT — CLEANING THE SCRIPT GRAVEYARD

Ghost IaC debt (documented in NoOps S3): thousands of cloud resources instantiated by abandoned Terraform scripts, active and billing, with no owner. A quarterly IaC audit lists all cloud resources, associates them with an owner and active use, and deletes orphans. Tool: terraform state list + AWS Cost Explorer. First audit duration: one day. Typical savings: 15 to 30% of monthly cloud invoice in organisations without IaC discipline.

◆ CLOUD-FREE BCP — THE QUESTION NOBODY ASKS

Every critical organisation must answer this question: 'What happens if AWS is unavailable for 24 hours?' If the answer is 'we don't know' — it is a security emergency, not a future project. The cloud-free Business Continuity Plan (BCP) defines critical workloads, data that must survive, and the failover path to backup infrastructure (bare-metal, sovereign colocation). Drafting a minimal BCP: one week with a competent infrastructure engineer. Cost: the engineer's cost, no additional infrastructure.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE HARDWARE EXISTS — AND IT IS CHEAPER THAN BELIEVED

◆ ⚡ IMMEDIATE — GRADE A REFURBISHED AND TPM CONTRACTS

Grade A Refurbished (documented in The Strategic Refurbished, 4p): servers 3 to 5 years old, grade A (functional, cleaned, tested), with chips already fabricated by TSMC — independent of future production. Price: a Dell PowerEdge R630 (2 Xeon E5 CPUs, 128 GB RAM, 4 SSD drives): €800 to €1,500 refurbished. Same capacity on AWS EC2: €8,000 to €15,000/year. The financial break is documented.

TPM contracts (Third Party Maintenance): physical maintenance was the lock the corpus had not resolved. Specialised actors (Park Place Technologies, Curvature, sovereign regional actors) ensure hardware maintenance of out-of-warranty servers at documented costs 40 to 60% below manufacturer contracts (Dell ProSupport, HPE Care Pack). The Accompanied Exodus systematically integrates a TPM contract into the TCO calculation. A line in the synthesis table — not an option.

◆ 🔧 MEDIUM TERM — THE ACCOMPANIED EXODUS: PORTABLE WORKLOADS FIRST

The Accompanied Exodus does not migrate everything overnight. It identifies portable workloads (Portability Score 2 or 3) and starts with them:

Phase 1 (0-6 months): Development and test environments → local bare-metal. No risk, immediate savings, team training without production pressure.

Phase 2 (6-18 months): Staging workloads and cold data → sovereign colocation. Reduced cloud surface without impacting production.

Phase 3 (18-48 months): Progressive production → sovereign bare-metal for workloads with Portability Score 2+. Proprietary workloads (Score 0-1) remain on cloud pending planned rewrite.

The Sovereign Interface: during transition, a hybrid network allows bare-metal and cloud workloads to coexist. The goal is not purity — it is progressive mastery.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

QUANTIFIED PRECEDENTS — CLOUD EXIT IS DOCUMENTED AND PROFITABLE

◆ QUANTIFIED PRECEDENTS — DROPBOX AND BASECAMP

These numbers are public, verifiable, and devastating to the 'cloud is always cheaper' narrative:

Dropbox (2017): infrastructure migration from AWS to own datacentres. Economy documented in the IPO prospectus (S-1, 2018): $74.6 million saved over 2 years. The repatriation required an infrastructure engineering team — exactly the profile NoOps had declared obsolete.

Basecamp (2023): David Heinemeier Hansson, co-founder, publicly documented the cloud exit. Economy: over $1.5 million per year. The calculation: $3.2 million annual cloud invoices → $500,000 amortisable hardware over 5 years + $200,000 operational costs. Bare-metal profitability achieved in under a year for their volume.

These two examples are not nostalgic exceptions — they are industrial margin optimisation decisions documented in legal filings and public publications.

◆ 🏛️ LONG TERM — RISC-V · ARM · EUROPEAN CHIPS ACT

The long-term hardware horizon is documented in The Taiwan Bottleneck (6p). TSMC fabricates ~90% of advanced chips in the Taiwan Strait. The long-term sovereign answer:

RISC-V: open instruction set architecture, royalty-free, already used by SiFive, Alibaba, Western Digital. RISC-V chips exist — not yet competitive with NVIDIA A100 for AI, but will be in 5 to 10 years for standard workloads.

European Chips Act (2023): €43 billion to double Europe's share of global semiconductor production by 2030. Sovereign fabs: TSMC Dresden (under construction), Intel Magdeburg (under construction), STMicroelectronics in France. The horizon is 10-15 years — but the work starts now.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

STRATUM 5 · LEGAL — ULTIMATUM · TERMS TRAINING · CLOUD ACT · SCHREMS II

THE LAW ALREADY EXISTS — IT MUST BE READ, TRAINED, AND USED

◆ ⚡ IMMEDIATE — THE ULTIMATUM OF 8 SEPTEMBER 2026 AND TERMS TRAINING

The ultimatum of 8 September 2026 makes four demands of AWS, GCP and Azure. These demands are the immediate contractual lever for any organisation renegotiating a cloud contract:

1 · Cancellable commits — without major financial penalty within 24h. Digital equivalent of a strategic veto right.

2 · Elimination of egress fees — or their legal capping (EU mobile roaming model).

3 · Guaranteed data portability — not in general Terms, in the signed contract.

4 · CLOUD Act position clarification — contractual commitment on non-disclosure without prior notification.

Legal training CIO/CFO: a 2-hour training on the CLOUD Act, Schrems II, and reading Terms §14.12 is the cheapest legal bifurcation. Can be delivered by any data law specialist. Cost: €500 to €2,000. Should be mandatory before signing any cloud contract exceeding €50,000 annually.

◆ 🔧 MEDIUM TERM — THE REVERSIBILITY CLAUSE AND STRICT SecNumCloud

The reversibility clause: any cloud contract exceeding 12 months must include a negotiated reversibility clause — exit plan, data export format, maximum timeline, capped cost. Not an extraordinary demand — it is what well-negotiated SaaS contracts already include for customer data.

Strict SecNumCloud: ANSSI's SecNumCloud qualification requires immunity from extraterritorial laws. In its strict version (without hyperscaler derogation), it is the only label guaranteeing CLOUD Act immunity for public administration data. Requiring it in sensitive data public tenders is not favouritism — it is a GDPR and sectoral regulatory compliance requirement.

LEX

Amine RAITI · CC BY-NC-SA 4.0

THREE LEGALLY UNASSAILABLE ELIMINATING CRITERIA — WITHOUT FAVOURITISM RISK

◆ THE PROBLEM — THE PROCUREMENT CODE AND THE FACE PRICE

The Public Procurement Code prohibits arbitrarily excluding actors under penalty of favouritism offence. Hyperscalers structured their offers to answer public tender requirements point by point — and the immediate face price criterion favours cloud (Free Tier, initial discounts) over Bare-Metal initial investment. Public buyers choose AWS not through favouritism — but through ease, cognitive availability (Kahneman's heuristic), and least-effort conformity (Default Effect).

◆ THE SOLUTION — THREE FUNCTIONAL AND LEGAL ELIMINATING CRITERIA

These criteria do not target AWS by name — they define objective functional requirements. Any entity that meets them can bid. Any entity that does not cannot — including AWS, Google and Microsoft as long as their European subsidiaries cannot guarantee CLOUD Act immunity.

Criterion 1 — Documentable extraterritorial immunity: the contract award entity must demonstrate it is subject to no law allowing a third-party State to access data without diplomatic procedure (MLAT). Legal basis: GDPR article 48 · Schrems II ruling (CJEU C-311/18, 2020) · DINUM Cloud au Centre Circular (2021).

Criterion 2 — 10-year TCO egress included: total cost of ownership must include exit costs (egress fees, migration costs, application rewrite). A cloud with €0 entry and €90,000 exit per PB is not cheaper than bare-metal at €50,000 initial investment. This criterion makes the financial comparison honest.

Criterion 3 — Contractually guaranteed data portability: data export in an open, documented format, within a defined maximum timeline, at capped cost, is a mandatory contractual clause.

LEX

Amine RAITI · CC BY-NC-SA 4.0

6

STRATUM 6 · INSTITUTIONAL — CPF · DINUM · ESN · PUBLIC SOVEREIGN CLOUD

CLOSING THE BUDGET LOOP — THE TAX THAT FUNDS ITS OWN LIBERATION

◆ ⚡ IMMEDIATE — THE BALANCED CPF AND ESN AS MARKET VECTORS

Balanced CPF: for every hyperscaler certification funded by the CPF, an equivalent neutral certification must be equally visible and fundable in the Mon Compte Formation catalogue. Not removing AWS Certified — restoring parity. Actor: France Compétences. Administrative, not legislative decision. Timeline: 3 to 6 months.

ESN as market vectors: ESNs (Capgemini, Sopra Steria, Atos, CGI) are cited in the corpus as hyperscaler partners in conflict of interest. But they also have another function: they train and employ tens of thousands of engineers in France. An ESN creating a 'Sovereign Bare-Metal Infrastructure' offer and training its teams accordingly is a massive market transformation lever. Not utopian — if public tenders require sovereignty (Stratum 5), ESNs will offer sovereign skills. ESNs follow markets. Stratum 5 creates the market. Stratum 6 trains the actors.

◆ 🔧 MEDIUM TERM — SOVEREIGN DINUM AND PUBLIC SOVEREIGN CLOUD

DINUM reorientation: the Cloud au Centre programme retained AWS, Azure and Google as public administration providers. Redirect part of the DINUM budget toward sovereign public infrastructures — OpenStack hosted on physical servers in France, operated by civil servants trained on neutral fundamentals.

The Public Sovereign Cloud: not a private company with a label (Bleu, S3NS) — a public infrastructure operated by the State, on the model of La Poste for mail or Caisse des Dépôts for public finance. Not statism — the answer to national critical infrastructure. Estimated budget: €500 million to €1 billion over 5 years. Savings on current administration cloud contracts: documentable from DINUM data.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

TRANSVERSAL STRATUM — CYBER INSURANCE AND THE NATIONAL TRAINING PLAN

◆ 🔧 MEDIUM TERM — SECNUMCLOUD AS INSURABLE STANDARD · FFA + ANSSI LOBBYING

Cyber insurance companies require hosting with certified third parties (often hyperscalers with their SOC 2 and ISO 27001) to cover operating loss risk. Uncertified bare-metal is perceived as higher risk — not because it objectively is, but because it lacks a standardised label recognised by insurers.

The answer: lobbying with the French Insurance Federation (FFA) and ANSSI for strict SecNumCloud to be recognised as equivalent to cloud ISO 27001 in cyber policies. Arguments:

1 · Cloud lock-in and Taiwan Bottleneck dependency constitute a systemic bankruptcy risk higher than locally mastered Bare-Metal infrastructure (documented in The Hand Over Nations, 8p).

2 · The Schrems II ruling creates documented legal risk for any organisation using American cloud services for GDPR data — this risk should increase non-compliant cloud premiums, not reduce them.

This lobbying is a medium-term lever (2-3 years). Actors: ANSSI + FFA + CNIL + DSI professional associations.

◆ 🏛️ LONG TERM — NATIONAL TRAINING PLAN AND EUROPEAN CHIPS ACT

National training plan: the equivalent of the 1960s Plan Calcul — adapted to the digital sovereignty era. Integrate into major engineering schools (Polytechnique, Centrale, INSA, IUT) a mandatory 'digital physics' module: network low-level layers, operating systems, hardware and Capacity Planning. Not cloud training — training in fundamentals that make any choice possible.

European Chips Act: €43 billion for European semiconductor manufacturing. TSMC Dresden · Intel Magdeburg · STMicroelectronics. The horizon is 2030-2035. The corpus documented the dependency (The Taiwan Bottleneck, 6p) — the Chips Act is the ongoing political response. France must ensure these investments produce chips usable by its own sovereign infrastructures, not only exported.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE COMPLETE ARCHITECTURE IN ONE TABLE — THE MAP AND ITS LEVERS

STRATUM	⚡ IMMEDIATE (0-6M)	🔧 MEDIUM TERM (6M-3Y)	🏛️ LONG TERM (3-10Y)	ACTORS
1 · School	€345 · Dish + RPi · 3D project	Official Anti-Amputation Core · 4h/term	National plan · Engineering school mandatory	MEN · *Teachers* · RNCP
2 · Training	56h Linux/Network/Storage · neutral CPF	Women 30+ · RNCP · ROME codes · Peer mentoring	Engineering school plan · National CPD	*France Compétences* · DARES · *Pôle Emploi*
3 · Organisation	Arch. question · Titles · IaC audit	Portability Score JIRA · Cap. Planning · BCP	Mature sovereign IDP · Peer mentoring	CIO · HR · *Infra Engineer* · PO
4 · Hardware	Grade A Refurbished + TPM contract · €800-1500	Accompanied Exodus · 3 phases · portable workloads	RISC-V · ARM · Sovereign fabs · Chips Act	CIO · *Infra Engineers* · *TPM providers*
5 · Legal	2h Terms training · Reversibility clause · Ultimatum	Sovereign specs · 3 eliminating criteria · 10-yr TCO	EU egress fee law · Insurable SecNumCloud	CIO · CFO · *Lawyers* · ANSSI · CNIL
6 · Institutional	Balanced CPF · ESN partners · DINUM orientation	Public Sovereign Cloud · FFA/ANSSI Lobbying	Chips Act €43B · National engineering plan	*France Compétences* · DINUM · FFA · *Ministries*

◆ NOTE — EACH STRATUM CAN BE ACTIVATED INDEPENDENTLY

Partial exit is preferable to paralysis. An organisation activating only Strata 3 and 4 progresses. One activating only Stratum 5 progresses legally. The effect is cumulative. All corpus references are available in the 64 preceding studies. cf. Anatomy of Digital Perdition (16p) for the captivity map · cf. The Architecture of Consent (8p) for cognitive biases · cf. The Hand Over Nations (8p) for the geopolitical framework.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

CLOSING · THE TEMPORAL COST OF BIFURCATION

THE BIFURCATION AT SCHOOL COSTS €345 AND A DISH — THE ONE AT T=20 YEARS COSTS A TOTAL REWRITE

The Operation Dindon corpus documented perdition in 64 studies. This study documents the exit in 6 strata. None is sufficient alone. Together, they form the only response proportionate to the architecture of captivity.

Stratum 1 (school) costs €345 and a satellite dish. Stratum 3 (Portability Score JIRA) costs 2 hours of configuration. Stratum 4 (Grade A Refurbished) costs €800 to €1,500 per server. Stratum 5 (legal training) costs €500 to €2,000. These interventions are available now, funded from existing budgets, without legislative decision.

Stratum 6 (Public Sovereign Cloud, Chips Act) requires political decisions and 5 to 15-year horizons. It cannot wait for Strata 1 to 5 to complete — it must be initiated now to be operational when the other strata have produced their effects.

HUMAN

Digital captivity was not built in a day.

It is not undone in a day.

But every lever activated is a link in the chain that gives.

And the chain gives.

Amine RAITI · Operation Dindon · 2026

NEMO SUPRA LEGEM EST

CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE SOVEREIGN INTERFACE

Inverted Abstraction as a Hyperscaler Detox Tool

Architectural Pattern for Technical Sovereignty

◆ **THE CENTRAL CONCEPT**

The Sovereign Interface is an abstraction layer whose contract is defined by the business domain — not by the vendor. It is *inverted* relative to classical adaptation: it is not the application that adapts to the vendor (BigQuery, DynamoDB, Pub/Sub), it is the vendor that adapts to the application. Changing vendor = changing the adapter. The application does not know. It is the architectural equivalent of the electronic adapter: the component that changes without either side changing.

RATIO

PATTERN	DOMAINS	WATERMARK
ANTI-LOCK-IN ACL	BIGQUERY DYNAMO · S3 PUB/SUB	RATIO

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · THE VOCABULARY LOCK-IN MECHANISM

BIGQUERY IN THE CODE IS MORE DANGEROUS THAN BIGQUERY IN THE INFRASTRUCTURE

Hyperscaler lock-in has two levels. The infrastructure level — VPC, subnets, load balancers, availability zones — is painful to migrate but technically feasible with Terraform and patience. The application level — BigQuery in the Python code, DynamoDB in the Java code, Pub/Sub in the Go code — is structurally different. It is the hyperscaler's dialect engraved in the business logic of the application. Migrating means rewriting. And rewriting a production application is a high-risk project with unpredictable cost and maximum regression probability.

◆ THE DIALECT AS AN INVISIBLE CHAIN

When a developer writes `client = bigquery.Client()`, they are not just choosing a service. They are committing the application to BigQuery's thinking schema — its data types, its extended SQL syntax, its ARRAY_AGG, STRUCT, proprietary partitioning and clustering. Six months later, these idioms are everywhere in the code. The application no longer speaks the business domain language — it speaks BigQuery.

The same dynamic applies to DynamoDB (key-value data model with GSI/LSI secondary indexes, proprietary condition expressions), to Pub/Sub (push/pull subscription model, acknowledgment, retry policy), to Cosmos DB (SQL, MongoDB, Cassandra API depending on the day — but all proprietary in their subtleties). Each hyperscaler service is a dialect. And a developer fluent in BigQuery is not fluent in PostgreSQL — they must learn a new dialect with its own traps, performance characteristics and limits.

◆ EGRESS FEES ARE THE SYMPTOM — THE DIALECT IS THE DISEASE

The Opération Dindon corpus documented egress fees as an economic capture mechanism. But even if egress fees were zero, migration would remain costly — because the true cost of migration is not data transfer, it is rewriting application code. Zero egress fees on S3 change nothing if the application calls 847 different S3 API endpoints with their nuances of versioning, presigned URLs and event notification. The gentle exit from hyperscaler cloud goes through code architecture, not just the network.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE ELECTRONIC ANALOGY — THE ADAPTER THAT LIBERATES

RS232 → USB · 5V → 3.3V · THE COMPONENT THAT CHANGES WITHOUT EITHER SIDE CHANGING

In electronics, when two incompatible systems must communicate, an adapter is placed between them. A 5V microcontroller talks to a 3.3V sensor via a level shifter. An RS232 port talks to a USB computer via an FTDI converter. The adapter is disposable — if the sensor changes, change the level shifter, not the microcontroller. If the port changes, change the converter, not the computer. Both sides remain stable. Only the adaptation component changes.

◆ THE ANTI-CORRUPTION LAYER — THE DOMAIN-DRIVEN DESIGN PATTERN

Eric Evans, in Domain-Driven Design (2003), names this pattern the Anti-Corruption Layer (ACL). When a system must integrate an external system with a different domain model, the ACL translates between the two — without letting the external model contaminate the internal model. The application only sees its own model. The ACL handles translation to the external. Changing the external = rewriting only the ACL.

Applied to cloud services: the application only sees its domain interface (`DataWarehouse`, `MessageQueue`, `ObjectStore`). The ACL translates to BigQuery, DynamoDB, S3. Migrating to PostgreSQL, Kafka, MinIO = rewriting the ACL. The business application does not change.

◆ WHAT ALREADY EXISTS — AND WHAT IS MISSING

Terraform: infrastructure-as-code abstraction. Same HCL for AWS, GCP, Azure. Limited to infra — not proprietary managed services.

Apache Beam: data pipeline abstraction. Same pipeline on Dataflow (GCP), Flink, Spark. Successful abstraction — but specific to batch/stream pipelines.

Kubernetes: container runtime abstraction. A K8s pod runs on EKS, GKE, AKS or bare-metal without modification.

S3-compatible APIs (MinIO, Ceph, Cloudflare R2): application speaks S3, runs on any compatible backend. The success model to duplicate.

JDBC/ODBC: database access abstraction since the 1990s. Application speaks standard SQL, driver translates to DBMS dialect.

What is missing: abstraction of proprietary managed services — BigQuery, Pub/Sub, DynamoDB, Cosmos DB, Kinesis. No common standard. This is where lock-in is deepest and the Sovereign Interface most necessary.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE THREE IMPLEMENTATION PATTERNS

DOMAIN INTERFACE · TRANSLATION FACADE · NEUTRAL DSL

◆ PATTERN 1 — THE DOMAIN INTERFACE (THE CLEANEST)

The application defines its needs in business terms. Behind the interface, concrete implementations for each vendor.

```
# Neutral interface defined by the business domain
class DataWarehouse(ABC):
    def query(self, sql: str) -> DataFrame: ...
    def insert(self, table: str, data: dict) -> None: ...

# BigQuery adapter
class BigQueryWarehouse(DataWarehouse):
    def query(self, sql):
        return self.client.query(sql).to_dataframe()

# PostgreSQL adapter (migration)
class PostgreSQLWarehouse(DataWarehouse):
    def query(self, sql):
        return pd.read_sql(sql, self.engine)

# The application never changes
warehouse: DataWarehouse = BigQueryWarehouse()  # or PostgreSQLWarehouse()
result = warehouse.query("SELECT user_id, revenue FROM sales WHERE date > '2026-01-01'")
```

Changing vendor = changing one configuration line. The business application does not know.

◆ PATTERN 2 — THE TRANSLATION FACADE (MIGRATION WITHOUT REWRITE)

An HTTP proxy that receives proprietary API calls and translates them to an open source backend. The application keeps calling the DynamoDB API — but the proxy redirects to ScyllaDB or Cassandra.

```
# Config: DYNAMO_ENDPOINT=http://local-proxy:8080
# Application code does not change
dynamodb = boto3.resource('dynamodb', endpoint_url=os.getenv('DYNAMO_ENDPOINT'))
table = dynamodb.Table('users')
table.put_item(Item={'user_id': '123', 'name': 'Amine'})
# In prod: proxy → DynamoDB AWS
# In migration: proxy → ScyllaDB open source
# Zero application changes
```

Existing projects: DynamoDB Local (testing), Localstack (full AWS emulation locally), dynamo-cassandra-proxy (Apache).

◆ PATTERN 3 — THE NEUTRAL DSL (THE MOST AMBITIOUS)

A neutral query language above proprietary dialects. The application speaks a common DSL, the translation layer generates backend-specific code.

```
# Neutral DSL
result = query.from_table("sales")
             .select(["user_id", "revenue"])
             .where("date", ">", "2026-01-01")
             .limit(1000)
             .execute()  # translates to BigQuery, PostgreSQL, DuckDB...
```

Existing projects in this spirit: SQLAlchemy (neutral ORM), dbt (neutral transformation), Ibis (pandas-like API on BigQuery, DuckDB, Snowflake).

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · CONCRETE EXAMPLES — BIGQUERY, DYNAMODB, PUB/SUB, S3

FOR EACH HYPERSCALER SERVICE — THE NEUTRAL INTERFACE AND THE ADAPTER

◆ BIGQUERY → DataWarehouse INTERFACE

Specific lock-in: extended SQL (ARRAY_AGG, STRUCT, UNNEST), partition/clustering, JavaScript UDFs, built-in ML.PREDICT.

Neutral interface: `DataWarehouse.query(sql: str) → DataFrame` — standard SQL only, no proprietary extensions.

Open source alternative: DuckDB (in-process analytics, same performance on small volumes), PostgreSQL + TimescaleDB (time series), ClickHouse (open source columnar analytics).

Bridge tool: Ibis — same Python API on BigQuery, DuckDB, Snowflake, PostgreSQL.

◆ DYNAMODB → KeyValueStore INTERFACE

Specific lock-in: GSI/LSI model (global/local secondary indexes), proprietary condition expressions (`attribute_exists`, `begins_with`), DynamoDB Streams.

Neutral interface: `KeyValueStore.get(key) / put(key, value) / query(pk, sk_prefix)`.

Open source alternative: ScyllaDB (DynamoDB API compatible), Cassandra, Redis with modules.

Translation facade: dynamo-cassandra-proxy (Apache) — DynamoDB API as facade, Cassandra as backend. Zero application rewrite.

◆ PUB/SUB → MessageQueue INTERFACE

Specific lock-in: topic/subscription model, push/pull, acknowledgment deadline, ordering keys, dead letter topics.

Neutral interface: `MessageQueue.publish(topic, message) / subscribe(topic, handler)`.

Open source alternative: Apache Kafka, RabbitMQ, NATS. Standard: CloudEvents (CNCF) for message format.

Key tool: CloudEvents + one adapter per broker. Message format is neutral — only transport changes.

◆ S3 → THE SUCCESS MODEL TO DUPLICATE

S3 is the only hyperscaler service whose API has become a de facto standard. MinIO, Ceph, Cloudflare R2, Backblaze B2 all implement the S3 API. The application speaks S3 and can run on any compatible backend. This is the model to replicate for other services. The lesson: standardise the API, not the service. The Sovereign Interface for S3 already exists — and it has freed thousands of organisations from object storage lock-in.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE CALCULATION THAT DOES NOT LIE

◆ COST OF IMPLEMENTING A SOVEREIGN INTERFACE

For a service like BigQuery used in a medium-sized Python application:

Define the domain interface: 2h — identify the operations actually used, reduce them to a minimal contract.

Write the BigQuery adapter: 4h — wrap existing calls behind the interface. Often this is refactoring existing code, not writing from scratch.

Write contract tests: 3h — tests that verify any adapter respects the interface contract.

Write a second adapter (local DuckDB): 3h — proves the interface is truly neutral and accelerates local development (no GCP needed in dev).

Total: 12h of senior engineering — a day and a half. Cost: ~€1,200-2,400 depending on rate.

◆ COST OF MIGRATION WITHOUT SOVEREIGN INTERFACE

For the same Python application, BigQuery → PostgreSQL migration without abstraction layer:

Code audit: 3 days — identify all BigQuery calls, all extended SQL idioms, all proprietary functions used.

Code rewrite: 10-30 days depending on complexity — replace each call, adapt each SQL query, handle type and behaviour differences.

Regression testing: 5-10 days — ensure the rewrite breaks nothing in business logic.

Progressive deployment: 3-5 days — switchover with rollback possible.

Total: 21 to 48 days of engineering. Cost: €25,000 to €60,000. And this assumes the migration succeeds on the first attempt — which is rare.

◆ THE ROI OF THE SOVEREIGN INTERFACE

12h initial investment → potential saving of 21 to 48 days during a migration. Ratio: 1 to 40 or 96. But the true return on investment is not in the migration — it is in the freedom to negotiate. An organisation whose application is decoupled from BigQuery can tell Google: "We are considering migrating to PostgreSQL in 6 months if pricing conditions do not change." This negotiating freedom has a value the calculation does not capture.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · TECHNICAL PRIMACY AS IMPLEMENTATION CONDITION

WITHOUT THE FINAL WORD — THE SOVEREIGN INTERFACE DOES NOT EXIST

The Sovereign Interface is an architectural decision. It cannot be imposed by a developer alone. It requires that the Principal SRE or Lead Architect has the final word on architecture choices — what the corpus calls Technical Primacy. Without it, the Sovereign Interface is a good idea that dies at the first sprint planning.

◆ THE SCENARIO WITHOUT TECHNICAL PRIMACY

Sprint planning. The backend developer proposes calling BigQuery directly — "it's simpler, we already have the SDK, it will work in two hours". The Principal SRE says "we need an abstraction layer — 12h of work but we'll be free to migrate". The product owner says "we don't have 12h, we have a Friday deadline". Without Technical Primacy, the product owner wins. BigQuery is in the code. The chain is in the code. And in 18 months, when Google raises prices by 40%, the migration will cost 48 days.

This is exactly Scenario 1 of "Technical Primacy" — the suboptimal decision made to meet a deadline, whose debt materialises 18 months later. With the Sovereign Interface, the deadline would have been 12h longer. Without it, it will be 48 days.

◆ THE SOVEREIGN INTERFACE AS TECHNICAL PRIMACY ARGUMENT

The Principal SRE who wants to impose an abstraction layer now has a quantified argument: "12h now, or 48 days in 18 months. Choose." This is no longer a discussion about abstract architectural best practices — it is an ROI calculation presentable to the board. Technical Primacy is easier to exercise when it relies on defensible numbers rather than engineering principles.

◆ THE DOCUMENTED VETO — APPLIED TO ARCHITECTURE

If management decides to override the Sovereign Interface recommendation, the Principal SRE documents their veto: "I advise against calling BigQuery directly without an abstraction layer. Estimated risk: forced migration within 2 years, estimated cost €25,000 to €60,000. Recommended alternative: Sovereign Interface, 12h, cost €1,200-2,400." This document protects the engineer. It records the decision. And when the migration arrives, the causal chain is traced.

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE SOVEREIGN INTERFACE IN THE CORPUS

THE FIFTH LAYER — ABSTRACT TO SURVIVE

The Opération Dindon corpus model had four layers: naming the body (The Ticket and the Talent), training it (The Foundation of Iron), making it visible (The Uniform of the Body), protecting it (The Digital Craftsmen). This study adds the fifth layer — abstract. Abstract application code from the hyperscaler dialect so that technical sovereignty is real, not merely declared.

◆ NO SOVEREIGNTY WITHOUT ABSTRACTION

The thesis "No Sovereignty Without Matter" applies to code: no sovereignty without abstraction. An organisation that declares its sovereignty but whose code speaks BigQuery is not sovereign — it is dependent with a sovereignty narrative. The Sovereign Interface is the matter of application sovereignty.

◆ THE TECHNICAL GENTLE EXIT

The Gentle Exit documented for DevOps engineers applies to architecture: migrate progressively, service by service, without a catastrophic D-Day. Plug in a new adapter. Test. Switch over. The application does not know. This is the gentle exit from hyperscaler cloud.

◆ THE FINOPS OF THE LAST GRAM

Abstraction allows routing to the cheapest backend by context. BigQuery for large analytical volumes in production. DuckDB locally for development (zero cost). PostgreSQL for intermediate volumes. The application does not change — only the adapter changes per environment.

◆ LAYER 5 — ABSTRACT

The Sovereign Interface
Anti-Corruption Layer
Neutral domain interface
→ Decouple to survive

◆◆◆

12h to implement. 48 days to migrate without it.
The engineer who places an abstraction layer today
is the engineer who negotiates from a position of strength tomorrow.
The metal precedes the code. The abstraction protects the code.

◆◆◆

NEMO SUPRA LEGEM EST

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE IT DIRECTOR'S DEPENDENCY AUDIT

Measuring dependency before deciding to leave

◆ DOCUMENTED FRENCH CASE STUDY

This study documents a real case observed on the French market, under French/EU law. Figures are in euros; legal references are to French and EU statutes. The mechanism it illustrates is recognised as universal across jurisdictions — the specific numbers, legal citations, and institutional frameworks (CSE, OPCO, RNCP, ANSSI, DGCCRF, etc.) are French-specific and are preserved here as documentary evidence, not as an action guide for another jurisdiction.

◆ WHO THIS IS FOR

For the exclusive use of IT departments and infrastructure architects. AWS · Google Cloud · Microsoft Azure. French market · law in force June 2026.

◆ SUMMARY

Page 2 — the 3 layers of dependency. Page 3 — the per-service evaluation grid. Page 4 — the 4 technical exit strategies. Page 5 — 15 pre-renewal questions.

AUDIT

Amine RAITI · CC BY-NC-SA 4.0

THE 3 LAYERS OF DEPENDENCY — MAP BEFORE DECIDING

◆ LAYER 1 — CONTRACTUAL DEPENDENCY (LEGAL TEAM'S REMIT)

Reserved Instances / Savings Plans: remaining term in months × monthly rate = total financial exposure. Extract from the billing console (AWS Cost Explorer · GCP Billing · Azure Cost Management). Azure Capacity Blocks: non-cancellable, non-refundable — check Product Terms (EES or MOSA) and active orders. AWS Bedrock throughput commitments: non-cancellable even after agreement termination — check the AWS Bedrock console under "Provisioned Throughput". These commitments are the incompressible financial floor of an exit. A progressive migration lets them run to natural term while the rest is migrated; early termination (possible under certain legal grounds) requires legal involvement.

◆ LAYER 2 — PRICING DEPENDENCY (EGRESS FEES)

Step 1: identify total data volume stored per service (S3/GCS/Azure Blob, databases, data warehouses, logs). Step 2: apply the provider's internet egress rate (AWS \$0.08–0.09/GB · GCP \$0.08–0.12/GB · Azure \$0.07–0.09/GB). For 100TB: €7,000–12,000 depending on provider and region. Step 3: request a written simulation of the egress cost for your exact volume — this written figure is enforceable in negotiation under Data Act art. 25. Note: some internet transfers are partially exempt (Azure: 5GB/month free · AWS: transfers to certain AWS services free) — calculate the net cost after exemptions before presenting the figure to management.

◆ LAYER 3 — TECHNICAL DEPENDENCY (PROPRIETARY APIS)

Level 1 — Low: open-standard service (S3-compatible storage · managed PostgreSQL · Kubernetes). Migration possible within weeks. Level 2 — Medium: proprietary API with a viable open-source alternative (Lambda → OpenFaaS · BigQuery → ClickHouse · Azure Functions → Knative). Partial refactoring: weeks to months. Level 3 — High: service with no direct equivalent (AWS Bedrock · Google Vertex AI · Azure OpenAI · certain managed ML services). Requires substantial rewrite or a change of approach: several months.

AUDIT

Amine RAITI · CC BY-NC-SA 4.0

THE EVALUATION GRID — SERVICE BY SERVICE

◆ HOW TO USE IT

Fill in this table for every cloud service used. A service scored Low/Low/Low is the first migration candidate. A service scored High/High/High should run to the end of its natural commitment while the rest is migrated. This is a prioritisation tool, not a migration decision.

Service / Use	Contractual dep.	Pricing dep. (egress)	Technical dep. (API)
Object storage (S3 · GCS · Blob)	MEDIUM — often under Savings	HIGH — volume = egress cost	LOW — S3-compatible standards
Managed DB (RDS · Cloud SQL · Azure SQL)	MEDIUM — often cancellable	MEDIUM — export possible	LOW — standard PostgreSQL/ MySQL
Serverless compute (Lambda · Functions)	LOW — pay-as-you-go	LOW — little data	MEDIUM — refactoring required
Data warehouse (BigQuery · Redshift · Synapse)	MEDIUM — committed use	HIGH — large volume	HIGH — proprietary SQL
Generative AI (Bedrock · Vertex · Azure OpenAI)	HIGH — non-cancellable throughput	LOW — little egress	HIGH — API with no equivalent
Managed Kubernetes (EKS · GKE · AKS)	LOW — cancellable	MEDIUM — application data	LOW — standard K8s
Identity & IAM (Cognito · Identity Platform · Entra)	LOW — cancellable	LOW — light data	HIGH — deep integration
CDN & network (CloudFront · Cloud CDN · Azure CDN)	LOW — cancellable	HIGH — high CDN egress	LOW — HTTP standards

◆ MIGRATION PRIORITIES

LOW/LOW/LOW → migrate first, minimal impact. MEDIUM → plan over 3–6 months. HIGH → let run to natural term.

◆ PRACTICAL RULE

Start by migrating what costs nothing to exit. This reduces the overall bill and creates a useful operational precedent for more complex migrations.

AUDIT

Amine RAITI · CC BY-NC-SA 4.0

THE 4 TECHNICAL EXIT STRATEGIES

◆ STRATEGY 1 — LIFT AND SHIFT · FAST, MODERATE COST

Reproduce the existing architecture at another cloud provider without application rewrite. Suited to containerised (Kubernetes) or classic IaaS workloads. Advantages: short timeframe (4–12 weeks), low operational risk, controlled migration cost. Limits: does not reduce dependency on proprietary managed services — lock-in transfers to the new provider if equivalent services are chosen there. When to use: IaaS or containerised workloads with no strong managed-service dependency, limited migration budget, short exit deadline imposed by management.

◆ STRATEGY 2 — PROGRESSIVE REFACTORING · LONG, HIGH ROI

Replace proprietary APIs with open or multi-cloud alternatives (Lambda → OpenFaaS/Knative · BigQuery → ClickHouse/DuckDB · Bedrock → self-hosted open-source models). Advantages: structurally reduces technical lock-in, improves long-term portability, each migrated service is an additional negotiating lever. Limits: long timeframe (6–18 months depending on complexity), high development cost, requires team upskilling. When to use: organisation with strong technical teams, cloud spend above €500k/year, long-term strategic decision to reduce dependency.

◆ STRATEGY 3 — MULTI-CLOUD · NEGOTIATING LEVER

Spread workloads across several providers to reduce reliance on one and create a real negotiating lever. Even a partial presence with a competitor makes the threat of migration credible. Advantages: immediate negotiating lever, operational continuity assured, reduced concentration risk. Limits: increased operational complexity, training and tooling costs, does not remove existing non-cancellable commitments.

◆ STRATEGY 4 — PARTIAL ON-PREMISE RETURN · COMPLIANCE, SENSITIVE DATA

Repatriate sensitive data and workloads to internal infrastructure hosted in France, in line with the "Cloud au centre" directive for public entities. Advantages: eliminates the US CLOUD Act issue, full compliance with the "Cloud au centre" doctrine, complete data control. When to use: health data · tax data · sensitive personal data · public entities subject to the "Cloud au centre" doctrine.

AUDIT

Amine RAITI · CC BY-NC-SA 4.0

15 QUESTIONS TO ASK BEFORE SIGNING OR RENEWING

◆ TECHNICAL QUESTIONS

Which services use non-portable proprietary APIs? (Bad answer: "all our services are standard" — ask for a written exhaustive list.) What is the estimated egress cost on our current data volume? What format are our data exportable in? (Proprietary format = technical dependency; demand a documented open format.) What migration tools do you provide if we decide to leave? (No tool = an obstacle to switching under Data Act art. 25.) Are your SLAs contractually guaranteed or "best effort"? What are the notice periods before each service is discontinued? Where are our data and encryption keys physically hosted? (Keys managed by the provider in the US = CLOUD Act exposure.)

◆ CONTRACTUAL QUESTIONS

Which language version of the ToS prevails? ("The English version prevails" = Toubon Law art. 5 violated — demand written confirmation that the French version is contractually binding.) Which commitments cannot be cancelled if strategy changes? Under what conditions can the Product Terms be modified? ("At any time without notice" = Civil Code art. 1193 violated.) Are egress fees capped or negotiable in our contract? Which court has jurisdiction in case of dispute? Have you signed the Data Act Addendum? (AWS has incorporated it, s.1.28 — ask GCP and Azure for the equivalent.) Are ingestion fees guaranteed at zero for the contract term? Is your AI opt-out (disabling model training on our data) activated and confirmed in writing?

◆ SOURCE

Based on legal analyses of AWS · GCP · Azure ToS, MOSA & EES, June 2026. Per aspera ad astra.

AUDIT

Amine RAITI · CC BY-NC-SA 4.0

THE REAL COST OF CAPTIVITY

What your cloud contract really costs you

◆ DOCUMENTED FRENCH CASE STUDY

This study documents a real case observed on the French market, under French/EU law. Figures are in euros; legal references are to French and EU statutes. The mechanism it illustrates is recognised as universal across jurisdictions — the specific numbers, legal citations, and institutional frameworks (CSE, OPCO, RNCP, ANSSI, DGCCRF, etc.) are French-specific and are preserved here as documentary evidence, not as an action guide for another jurisdiction.

◆ WHO THIS IS FOR

For Finance Departments and General Management. AWS · Google Cloud · Microsoft Azure. French market, June 2026.

◆ SUMMARY

Page 2 — the real cost of captivity. Page 3 — stay vs leave vs multi-cloud matrix. Page 4 — budgeting an exit. Page 5 — the argument for management.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE REAL COST OF CAPTIVITY

◆ COST 1 — EGRESS FEES: THE PRICING PERIMETER WALL

Egress fees do not appear in any ToS clause — they are in a separately published price list. A decision-maker signing a cloud contract generally does not see this amount.

Data volume	AWS	GCP	Azure
10 TB	€800–900	€800–1,200	€700–900
100 TB	€8,000–9,000	€8,000–12,000	€7,000–9,000
1 PB (large org.)	€80–90k	€80–120k	€70–90k

◆ COST 2 — NON-CANCELLABLE COMMITMENTS: THE INVISIBLE FINANCIAL EXPOSURE

Formula: exposure = Σ (remaining months × monthly rate) for each non-cancellable commitment. AWS Reserved Instances · AWS/Azure Savings Plans · Azure Capacity Blocks · AWS Bedrock throughput · GCP Committed Use Discounts — none of these are refunded on early termination (GCP s.8.8: no refund · Azure EES: 12% fee + $50k/yr cap). Example: 100 AWS Reserved Instances at $2,000/month each over 3 years with 18 months remaining = $3,600,000 of incompressible financial exposure if terminated immediately. This figure appears in no standard dashboard.

◆ COSTS 3, 4, 5 — MIGRATION · TRAINING · DOUBLE-RUN

Technical migration cost: 10–30% of annual cloud spend depending on complexity — includes refactoring of proprietary APIs, testing, data validation. Training cost: team upskilling on the new platform, certifications, transitional productivity loss. Double-run cost: the period (typically 3–6 months) during which both environments coexist — non-cancellable commitments continue to be billed throughout.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE DECISION MATRIX — STAY · LEAVE · MULTI-CLOUD

Criterion	Stay	Leave	Multi-cloud
3-year total cost	Known but captive	High short-term	Medium, complex
Contractual risk	High — mutable ToS	Reduced post-migration	Shared across 2 providers
Strategic flexibility	Low — dependent	High — full freedom	Medium — improved
Regulatory compliance	Exposed — CLOUD Act	Controlled if sovereign	Partial
Operational effort	Low — status quo	High — 6–18 months	Medium — ongoing

◆ SME — under €100k/year cloud spend

Recommendation: multi-cloud or partial exit. At this level, non-cancellable commitments are generally low. The main lever is contractual negotiation (void clauses) plus refusing to renew Savings Plans.

◆ MID-CAP — €100k–1M/year cloud spend

Recommendation: audit first, decide second. Financial exposure can be significant. Map non-cancellable commitments before any decision. Involve legal to invoke void clauses as a renegotiation lever.

◆ LARGE ORGANISATION — over €1M/year cloud spend

Recommendation: multi-year migration programme sequenced by layer. Exit cost can be high (egress + commitments + migration) but regulatory exposure is equally high (CLOUD Act · Data Act · Directive 2018/1972). Migration ROI depends directly on the length of remaining commitments — a contractual audit is mandatory before deciding. Engage a specialised lawyer and technical support for migration costs above €500k.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

BUDGETING AN EXIT — THE OFTEN-FORGOTTEN LINE ITEMS

◆ WARNING ON ESTIMATES

The ranges below are indicative, based on general sectoral observations. They vary significantly with organisation size, architectural complexity, level of proprietary API dependency, and chosen migration strategy. These figures are meant to frame a projected budget — not to serve as a quote.

Cost item	Basis for calculation	Often forgotten?
Data egress fees	Data volume (GB) × provider rate	YES — outside ToS
Remaining non-cancellable commitments	Remaining term × monthly rate	YES — often ignored
Proprietary API refactoring	10–30% of annual cloud cost	YES — underestimated
Double-run period (3–6 months)	Current cloud cost × double-run duration	YES — systematic
Team training on new platform	Nb people × certification cost × nb certifications	SOMETIMES
Post-migration validation and testing	10–20% of technical migration cost	SOMETIMES
Legal support	€2,000–10,000 depending on complexity	NO — usually budgeted

◆ WHEN IS EXIT FINANCIALLY WORTHWHILE?

Migration is financially worthwhile when the total 3-year cost of captivity (cloud bill + contractual risk + regulatory exposure) exceeds the total migration cost. This comparison must include the cost of regulatory inaction (CLOUD Act · Data Act) and the risk of unilateral price increases during non-cancellable commitments. For an organisation with cloud spend above €500k/year, a formal profitability analysis is warranted before any decision.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE ARGUMENT FOR MANAGEMENT

◆ ANGLE 1 — FINANCIAL RISK

"Our exposure to non-cancellable commitments represents X million euros. If our strategy changes, if revenue falls, if we are acquired, these commitments continue to be billed. This risk is not in our dashboard." → Present the total financial exposure calculation (Fiche B, Layer 1).

◆ ANGLE 2 — REGULATORY RISK

"Our data is subject to the US CLOUD Act. The ToS we signed place us under the jurisdiction of the State of Washington. The 2023 EU Data Act gives us rights our contracts do not yet incorporate." → Reference the Microsoft EES analysis: a €152m French Ministry of Education market under US law.

◆ ANGLE 3 — STRATEGIC RISK

"We cannot change provider without prohibitive cost. This means our provider can change prices in 6 months and our only recourse will be to pay or migrate in an emergency." → Present the unilateral price modification clause (AWS s.50.12.3 · GCP s.2.6) combined with commitment non-cancellability.

◆ 3 QUESTIONS FOR THE ACCOUNT MANAGER AT RENEWAL — ASK IN WRITING BEFOREHAND

1. "Can you confirm in writing that the French version of your ToS is contractually binding and that the English version does not prevail?" 2. "What will our total egress cost be if we migrate all our data within 12 months? Can you provide this estimate in writing?" 3. "If your Product Terms change in the next 12 months, under what conditions can we terminate current commitments without penalty?"

◆ WHY "WE'VE ALWAYS DONE IT THIS WAY" IS NO LONGER ENOUGH

The EU Data Act (2023/2854) is in force. It renders unenforceable clauses that obstruct a change of provider. The global market is evolving: AWS has incorporated a Data Act Addendum (s.1.28) and a UK Switching Addendum (s.1.30) under regulatory pressure. What was contractually acceptable in 2020 is contestable today. Inaction in the face of clauses identified as illegal is not a neutral position — it is a documented risk.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE 5 STEPS OF A SUCCESSFUL EXIT

From decision to effective termination

◆ **DOCUMENTED FRENCH CASE STUDY**

This study documents a real case observed on the French market, under French/EU law. Figures are in euros; legal references are to French and EU statutes. The mechanism it illustrates is recognised as universal across jurisdictions — the specific numbers, legal citations, and institutional frameworks (CSE, OPCO, RNCP, ANSSI, DGCCRF, etc.) are French-specific and are preserved here as documentary evidence, not as an action guide for another jurisdiction.

◆ WHO THIS IS FOR

Operational guide, all departments. AWS · Google Cloud · Microsoft Azure. French market, June 2026.

◆ SEQUENCE

1 Trigger the audit D-0 → 2 Contractual audit D+7 to D+21 → 3 Technical audit D+14 to D+45 → 4 Negotiation D+30 to D+60 → 5 Migration execution D+45 to D+180.

◆ SUMMARY — 6 PAGES

Page 2 — Steps 1 & 2: decide and audit. Page 3 — Steps 3 & 4: prepare and negotiate. Page 4 — Step 5: execute the exit. Page 5 — Vendor blocking scenario. Page 6 — Dashboard (steps 1–5).

OPERA

Amine RAITI · CC BY-NC-SA 4.0

STEPS 1 & 2 — DECIDE AND AUDIT

◆ STEP 1 — TRIGGER THE AUDIT — D-0 · Owner: General Management · IT

Who decides: the decision to audit does not yet mean a decision to migrate — it is a decision to gather information. Taken by top management or IT based on one or more triggers: falling revenue · ongoing acquisition · contract renewal within 6 months · a price increase announced by the provider · a new regulatory constraint. First information to gather: list of all active cloud contracts with expiry dates · list of all non-cancellable commitments with amounts and durations · latest applicable ToS (check whether these are Product Terms or a master agreement). What not to tell the provider yet: do not inform the account manager of an exit audit before completing Step 2 — premature disclosure can trigger sales calls that create confusion in internal decision-making. Estimated duration: 2–4 weeks to gather basic information. If contracts are not easily accessible internally, that is itself a signal of contractual risk.

◆ STEP 2 — CONTRACTUAL AUDIT — D+7 to D+21 · Owner: Legal

Identify all current commitments: expiry date · monthly amount · early-termination conditions · possible penalties (Azure: 12% + $50k/yr cap · AWS/GCP: no refund exception). Map clauses that are void as a matter of law (see Fiche A, page 3): English-version primacy · absolute non-cancellability · unilateral price modification without termination right · foreign jurisdiction — these clauses reduce real financial exposure compared with what the provider will show. Produce the internal legal memo for management and IT, listing real financial exposure (net of void clauses), applicable legal grounds, and available recourse if blocked. Expected outcome: management and IT know exactly what can legally be cancelled today, what cannot and until when, and what exit would cost in the best and worst cases.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

STEPS 3 & 4 — PREPARE AND NEGOTIATE

◆ STEP 3 — TECHNICAL & FINANCIAL AUDIT — D+14 TO D+45 · OWNER: IT · FINANCE

IT — map dependencies: use the Fiche B (page 3) grid to score each service, identify quick-to-migrate services vs. high-dependency ones, calculate egress cost on real data volume. IT — choose migration strategy: among the 4 strategies in Fiche B (page 4), identify the one matching the organisation's profile — this decision drives budget and timeline. Finance — model total exit cost: use the line items from Fiche C (page 4), add egress + remaining commitments + refactoring + double-run + training, compare with 3-year captivity cost (current bill × 3 + price-increase risk). Expected outcome: management has a costed stay-vs-leave comparison and a validated migration budget — the formal decision to migrate (or renegotiate) can be taken.

◆ STEP 4 — NEGOTIATION WITH THE PROVIDER — D+30 TO D+60 · OWNER: LEGAL · FINANCE · IT

When to open negotiation: after Step 2 (contractual audit complete) — never before identifying void clauses and legal grounds. Opening negotiation with an incomplete file exposes the organisation to unnecessary concessions. How to use void clauses as leverage: notify in writing that certain clauses breach the Toubon Law, Civil Code art. 1171, or Data Act art. 25. Request a written position from the provider — this is not a threat, it is a documented legal finding. What can be obtained without going to court: reduced egress fees for migration · extended termination notice without penalty · addition of a Data Act Addendum (AWS did this under regulatory pressure) · modification of governing-law clauses.

◆ 3 NEGOTIATION MISTAKES TO AVOID

Do not sign a "commercial" amendment without legal verification that it does not waive the rights being invoked. Do not accept consumption credit in place of a cash refund — that is additional pricing lock-in. Do not partially renew without first auditing which new clauses apply to the renewal.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

STEP 5 — EXECUTE THE EXIT · D+45 TO D+180

◆ RECOMMENDED SEQUENCE — IN THIS ORDER · OWNER: IT

Phase 1 — data first: export and validate data before anything else. Data loss during migration is the most critical, most irreversible risk — verify exported data integrity before deleting anything at the source provider. Phase 2 — stateless services next: microservices, APIs, data-processing services. Low operational risk, maximum reversibility — this phase demonstrates migration feasibility to the whole organisation. Phase 3 — stateful services: databases, message queues, identity services. Requires a synchronisation period between source and destination — a double-run delay is unavoidable. Phase 4 — non-cancellable services maintained: Savings Plans, Reserved Instances and other non-cancellable commitments continue to be billed during migration — keeping them active to their natural term reduces total exit cost.

◆ INTERNAL COMMUNICATION — OWNER: IT · HR

Cloud migrations affect technical teams who have built skills on the source platforms. Clear internal communication on the reasons for migration (strategic, financial, regulatory) and on the upskilling plan reduces resistance and turnover risk during the critical period. Migration documentation (architecture decisions, migration scripts, validation tests) is an organisational asset that must be produced and retained — it protects the organisation in any dispute with the source provider over the state of the data at exit.

◆ POST-MIGRATION VALIDATION — OWNER: IT · LEGAL

Before formally terminating source services: verify migrated data integrity · validate application performance on the new platform · confirm regulatory compliance (GDPR · Data Act · "Cloud au centre" for public entities). Retain migration evidence for 5 years (dated exports, validation logs, correspondence with the source provider) in case of later disputes over non-cancellable commitments or contested fees.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

BLOCKING SCENARIO — WHAT TO DO IF THE PROVIDER OBJECTS

◆ SITUATION A — The provider invokes a non-cancellable clause

Immediate response: notify in writing (email with read receipt) that the clause invoked breaches the EU Data Act (2023/2854) art. 25 (unenforceability of contractual obstacles to switching provider) and Civil Code art. 1171 (significant imbalance in an adhesion contract). Request written confirmation of the provider's legal position within 15 days. If the provider maintains its position: file with the DGCCRF (signal.conso.gouv.fr for professionals) with a documented file · contact a lawyer specialised in digital law for a formal notice · if a public entity: an administrative appeal (REP) may be possible before the Administrative Court (see Fiche A, page 4).

◆ SITUATION B — The provider applies prohibitive egress fees

Immediate response: notify that the egress fees constitute a pricing obstacle to switching provider under Data Act art. 25. Request an exemption or a significant reduction in the context of migration, referencing the portability obligations the regulation imposes on cloud service providers. AWS incorporated a Switching Addendum under regulatory pressure — use this precedent with GCP and Azure.

◆ SITUATION C — The provider suspends services during migration

Immediate response: document the suspension (screenshots, logs, correspondence) · notify in writing that the suspension constitutes non-performance under Civil Code art. 1217 · urgently engage a lawyer for an interim commercial injunction depending on the financial stakes.

◆ ABSOLUTE RULE IN A BLOCKING SITUATION

Every exchange with the provider must be documented in writing. A phone call with an account manager is not an enforceable position. Systematically demand written confirmation of the provider's position on every disputed point.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

EXIT DASHBOARD — MASTER CHECKLIST

◆ HOW TO USE

⚖ Legal · 🖥 IT · 💼 Finance · ★ Management. Use in steering-committee meetings.

Action	Owner	Timing
Identify the trigger(s) for the audit decision	★ Mgmt	D-0
Gather all active cloud contracts with expiry dates	⚖ Legal	D-0 to D+7
List all non-cancellable commitments (amounts + remaining terms)	⚖ Legal	D+7
Decide not to inform the provider at this stage	★ Mgmt	D-0
Identify clauses void as a matter of law (art.1171 · Toubon · Data Act art.25)	⚖ Legal	D+7 to D+14
Calculate real total financial exposure (with/without void clauses)	💼 Finance	D+14
Identify available recourse if blocked (DGCCRF · Data Act · Court)	⚖ Legal	D+14
Produce internal legal memo for management	⚖ Legal	D+21
Map technical dependencies by service (Fiche B grid)	🖥 IT	D+14 to D+30
Calculate egress cost on real data volume	🖥 IT	D+21
Choose migration strategy	🖥 IT	D+30
Model total exit cost · compare with 3-year captivity cost	💼 Finance	D+45
Prepare negotiation file (void clauses + Data Act art.25 grounds)	⚖ Legal	D+30
Send the 3 written questions to the account manager (Fiche C, page 5)	⚖ Legal	D+30
Document all written provider responses	⚖ Legal	Ongoing
Review any amendment proposed by the provider before signing	⚖ Legal	Before signing
Formal GO / NO GO migration decision by management	★ Mgmt	D+60
Export and validate data integrity (before any other action)	🖥 IT	Phase 1
Migrate stateless services · validate in production	🖥 IT	Phase 2
Migrate stateful services · double-run period · synchronisation	🖥 IT	Phase 3
Validate GDPR · Data Act · "Cloud au centre" compliance	⚖ Legal	Phase 4
Keep non-cancellable commitments to natural term — do not cancel early	💼 Finance	Throughout
Archive all migration evidence for 5 years	⚖ Legal	Final phase

OPERA

◆ SOURCE

Based on legal analyses of AWS · GCP · Azure ToS, MOSA & EES, June 2026. Summum ius, summa iniuria — Cicero. Nemo supra legem est.

Amine RAITI · CC BY-NC-SA 4.0

THE LOCK-IN OF SKILLS

HR · Change Management · Product Owners · Business Units

◆ DOCUMENTED FRENCH CASE STUDY

This study documents a real case observed on the French market, under French/EU law. Figures are in euros; legal references are to French and EU statutes. The mechanism it illustrates is recognised as universal across jurisdictions — the specific numbers, legal citations, and institutional frameworks (CSE, OPCO, RNCP, ANSSI, DGCCRF, etc.) are French-specific and are preserved here as documentary evidence, not as an action guide for another jurisdiction.

◆ CENTRAL THESIS

Skills lock-in is the most durable layer of vendor lock-in. A company that technically migrates to an open stack but whose teams still think in terms of "Lambda functions" and "BigQuery datasets" has not really migrated. And a business unit that expresses its needs in named products ("I want Azure OpenAI") rather than in features ("I want an LLM inference service") closes off competition before IT even begins to negotiate.

◆ SUMMARY

Page 2 — Part 1: Skills Lock-in, anatomy & profiles. Page 3 — Part 1: skills transformation plan. Page 4 — Part 2: Demand Lock-in, PO & business units. Page 5 — Synthesis: 6 layers, common 6-month action plan.

◆ SCOPE NOTE

This document is not legal advice nor individualised HR advice. For any specific situation, consult a firm specialised in digital transformation and change management.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

SKILLS LOCK-IN — ANATOMY & 5 RESISTANCE PROFILES

◆ HOW SKILLS LOCK-IN FORMS

AWS has 12 official certifications. GCP has 11. Azure has 9+. These programmes are free or heavily subsidised by the hyperscalers themselves — because an AWS-certified engineer is an engineer whose market value depends on AWS's continued use at their company. This is not malice, it is ecosystem design. Attachment bias runs both ways: the engineer does not want to devalue their certification, the company does not want to lose the engineer. Result: nobody proposes migrating.

◆ ① THE "FULL HYPERSCALER" ENGINEER — CERTIFIED, 3+ YEARS ON THE PLATFORM

Resistance: explicit fear of losing market value. "If we leave AWS, my certifications are worthless." Lever: show that CKAD, CKA, Linux Foundation, and OpenTofu certifications are more portable and better paid long-term — hyperscaler skills are a niche, open skills are a market. Note: this resistance is rational short-term — contextualise it over a 3–5 year career, don't deny it.

Note (national institutional term used in this study's HR context: CPF — Compte Personnel de Formation, France's individual training-rights account).

◆ ② THE "MANAGED-SERVICES COMFORT" ENGINEER — DEVOPS · SRE · OPS

Resistance: fear of extra operational work. "Before, I clicked in the console. Now I manage config files at 2am." Lever: show that modern IaC (Terraform/OpenTofu · Helm · ArgoCD) reduces manual work as much as the console does — the difference is not the amount of work but its nature. Note: hyperscaler consoles are UX masterpieces designed to create this dependency — naming it helps overcome it.

◆ ③ THE SENIOR "I BUILT ALL THIS" ENGINEER — 5+ YEARS ON THE EXISTING ARCHITECTURE

Resistance: identity attachment to the architecture. "I designed this system. Questioning it means questioning my decisions." Lever: bring them in as migration architect, not migration target — their knowledge of the existing system is irreplaceable, they become the project's most valuable expert. Note: they often have the best reasons to migrate — they know the technical debt nobody else sees.

◆ ④ THE TEAM MANAGER — TECH LEAD · HEAD OF ENGINEERING

Resistance: fear of delays and impact on OKRs. "If migration takes 6 months, I miss my Q3 goals." Lever: give realistic per-layer migration figures (see page 3 table), show that progressive migrations don't impact production, negotiate migration-specific OKRs. Note: this resistance often signals the project wasn't properly championed one level up — the problem is hierarchical, not technical.

◆ ⑤ MIDDLE MANAGEMENT — TECH LEAD · SCRUM MASTER · PROJECT MANAGER

Resistance: passive, invisible resistance — neither open opposition nor active support. Prioritises production bugs over migration tickets. "We'll see after the release." Lever: explicitly integrate migration into sprints with clear user stories and definitions of done, name a migration champion in each squad, make resistance visible and addressable. Note: the most dangerous profile because the least visible — unconscious sabotage can derail a migration project without ever appearing in a steering-committee report.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

SKILLS TRANSFORMATION PLAN — 3 PHASES

◆ MAPPING CURRENT SKILLS → REQUIRED SKILLS (ESSENTIAL PREREQUISITE)

Hyperscaler skill	Open equivalent	Gap	Training time
AWS Lambda / Cloud Functions / Azure Functions	OpenFaaS · Knative · Fission	Low	2–4 weeks
CloudWatch / Cloud Monitoring / Azure Monitor	OpenTelemetry + Prometheus + Grafana + Loki	Medium	3–6 weeks
BigQuery SQL / Redshift / Azure Synapse	ClickHouse SQL · DuckDB	Low	1–3 weeks
EKS / GKE / AKS	k3s · k0s · vanilla Kubernetes	Low	1–3 weeks
Terraform (BSL) + AWS/GCP/Azure console	OpenTofu + declarative IaC	Low	1–2 weeks
RDS / Cloud SQL / Azure SQL	Self-managed PostgreSQL + PgBouncer	Medium	3–6 weeks
SageMaker / Vertex AI / Azure ML	vLLM + Ray + MLflow + Hugging Face	High	2–3 months
S3 / GCS / Azure Blob (AWS SDK)	MinIO (identical S3 SDK) · Scaleway	None	0 — zero refactoring

◆ PHASE 1 · MONTH 1 — AWARENESS · LIFT ANXIETY, SPARK CURIOSITY

Discovery sessions on alternatives (not technical training) · sharing of successful migration experiences · live demos of ClickHouse, OpenTelemetry, k3s · no pressure on timelines. Goal: make alternatives real in teams' eyes, not convince them to migrate.

◆ PHASE 2 · MONTHS 2–3 — PILOT · PROVE IT ON A REAL, NON-CRITICAL CASE

One non-critical project migrated by a volunteer team (never imposed) · dual skills maintained during the period · feedback documented and shared with the whole team · volunteers become the natural internal trainers for phase 3. If partial failure: document why, not who.

◆ PHASE 3 · MONTHS 4–6 — ROLLOUT · GENERALISE WITHOUT BRUTALISING

Training in waves of 4–6 people · former hyperscaler engineers become internal trainers (valorisation, not devaluation) · open-standard certification valued in annual reviews · integration into job descriptions and evaluation criteria. Migration becomes an upskilling project, not a replacement.

◆ WHAT THE CHANGE MANAGER MUST NOT DO

✗ Present migration as a criticism of past decisions. ✗ Promise an unrealistic timeline ("we migrate in 2 weeks"). ✗ Train everyone simultaneously — create waves and ambassadors instead. ✗ Ignore middle management's passive resistance. ✗ Measure migration speed only — also measure quality of upskilling. ✗ Remove access to the old platform before the team is comfortable on the new one.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

DEMAND LOCK-IN — HOW BUSINESS UNITS CO-AUTHOR LOCK-IN

◆ IT WASN'T IT THAT CHOSE AZURE OPENAI. IT WAS MARKETING THAT WANTED AN AI COPILOT.

Demand lock-in works because business units express needs in terms of named products rather than features. "I want Azure OpenAI" instead of "I want an OpenAI-API-compatible LLM inference service". "I want BigQuery" instead of "I want a queryable OLAP data warehouse for large volumes". This semantic shift turns a functional need into a vendor commitment before IT even begins evaluating alternatives. IT is no longer the decision-maker — it executes a decision the business already made.

◆ REFORMULATION GUIDE — 10 EXAMPLES BY DEPARTMENT

To be systematised, in writing, before any IT request.

Department	❌ Named-product request	✅ Functional reformulation (opens competition)
🎯 Marketing	"We want Azure OpenAI for our customer chatbot"	"We want an LLM inference service · OpenAI-compatible API · latency <2s · 1,000 requests/day · customer data not used for third-party training"
📊 Data	"We want BigQuery for our data warehouse"	"We want an OLAP warehouse able to query 5TB · response <5s · 10 concurrent analysts · native CSV/Parquet export · standard SQL"
⚙️ Ops	"We want AWS Lambda for our event pipeline"	"We want an event-driven compute service · stateless · pay-per-invocation · Python 3.11 runtime · 15-min max timeout · Kafka integration"
📈 Finance	"We want Power BI integrated with Azure for our reports"	"We want a BI tool · SQL/CSV connectors · dashboard sharing by link · SSO authentication · PDF export · hosted within our data perimeter"
🛒 E-commerce	"We want DynamoDB for our product catalogue"	"We want a NoSQL document database · <10ms read latency · 10,000 req/s peak · flexible schema · native REST API · 99.99% SLA"
🤖 Product	"We want GitHub Copilot for our devs"	"We want a code-completion assistant · VS Code integration · configurable model · code data not sent for training · on-premise deployment possible"
👥 HR	"We want Microsoft Teams + Azure for collaboration"	"We want a collaboration suite · messaging + video + file sharing · <200 users · EU hosting · documented GDPR compliance · data export possible"
🏥 Compliance	"We want AWS HealthLake for our health data"	"We want a health-data warehouse · FHIR R4 format · France hosting · HDS-certified · REST API access · client-side key encryption · no non-EU subcontracting"
🔒 Security	"We want Microsoft Sentinel for our SIEM"	"We want a SIEM · 50GB/day log ingestion · customisable detection rules · export API · 90-day retention · SOAR integration · no single-cloud dependency"
📦 Logistics	"We want Google Maps Platform for geolocation"	"We want a geocoding + routing API · 1M requests/month · 99.9% SLA · route-data export · alternative: OpenStreetMap + Valhalla open-source"

◆ THE PRINCIPLE

Functional reformulation opens competition. A named-product request closes it before evaluation even begins.

Amine RAITI · CC BY-NC-SA 4.0

SYNTHESIS — THE 6 LOCK-IN LAYERS · COMMON ACTION PLAN

◆ 5 QUESTIONS EVERY PO MUST ASK BEFORE VALIDATING A STACK — IN WRITING, BEFORE ANY IT MEETING

1. "Which of our customers'/users' data will this service process?" → identify GDPR and CLOUD Act exposure before signing; if sensitive, require EU hosting and a French-law contract. 2. "Under which governing law is this service contracted?" → "State of Washington law" is a red flag. 3. "What is the exit cost if we change our mind in 18 months?" → force IT to quantify egress fees and non-cancellable commitments before signing. 4. "Is there an open-source or sovereign alternative covering 80% of the need?" → don't demand 100% parity; 80% functionality with 100% sovereignty is often the best trade-off. 5. "Can this service run in degraded mode if the provider is unavailable?" → a critical service with no fallback is absolute dependency, and the provider knows it.

◆ THE 6 LOCK-IN LAYERS — FULL ARCHITECTURE

Layer 1 Contractual: non-cancellable clauses, language clause, unilateral modification (Fiche A). Layer 2 Pricing: egress fees, termination charges, entry/exit asymmetry (Fiche C). Layer 3 Technical: proprietary APIs, managed services, superlinear effect (Fiche B). Layer 4 Skills: proprietary certifications, attachment bias, resistance to change (this fiche, p.2–3). Layer 5 Demand: named-product requests, prescribing business units (this fiche, p.4). Layer 6 Cognitive: KPIs defined in the provider's console, dashboard loss if migrating.

Period	HR · Change Manager	Business · PO	Indicator
Months 1–2	Certification mapping · Identify 5 profiles · Training plan	Audit named-product requests · Functional reformulation · 5 PO questions	Validated mapping report
Months 3–4	Pilot phase 2 · Volunteer team · Documented feedback	Validate alternative covers 80% of need · Go/No-Go service 1	1 service migrated, team trained
Months 5–6	Rollout phase 3 · Wave training · Open certifications valorised	Integrate PO questions into procurement · Train POs on demand lock-in	Review · next decision
Ongoing	Former hyperscaler engineers = internal trainers	Reformulation template in briefs · PO = first sovereignty filter	% requests reformulated

◆ SOURCE

This document is not individualised HR or legal advice. For organisations with 200+ employees, support from a firm specialised in digital transformation is recommended. Homo sum: humani nil a me alienum puto — Terence.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

SRE/OPS GUIDE · TECHNICAL DOCUMENT · GLOBAL MARKET · JUNE 2026

CLOUD COST REDUCTION & PROGRESSIVE MIGRATION

AWS · Google Cloud · Microsoft Azure

Operational guide for SRE, Ops and Infrastructure Architect teams
33 technical blocks · 3 cleaning sections · Reference open-source stack
Summary tables per hyperscaler

33 TECHNICAL BLOCKS	3 CLEANING SECTIONS	6 OPEN-SOURCE TABLES	3 SUMMARY TABLES

TABLE OF CONTENTS — 12 PAGES

- Page 2 — AWS · Compute & Network (5 blocks)
- Page 3 — AWS · Storage · DB · Serverless · Cleanup (5 blocks)
- Page 4 — AWS · Progressive Migration
- Page 5 — GCP · BigQuery · Compute · Network (5 blocks)
- Page 6 — GCP · Storage · DB · AI · Cleanup (5 blocks)
- Page 7 — GCP · Progressive Migration
- Page 8 — Azure · Compute · Network · Storage (5 blocks)
- Page 9 — Azure · DB · AI · Data · Cleanup (5 blocks)
- Page 10 — Azure · Progressive Migration
- Page 11 — Reference Open-Source Stack (synthesis)
- Page 12 — Quick Reference Tables AWS · GCP · Azure

This document is purely technical. It contains no legal references or political dimension.
Cost reduction estimates are indicative and based on common use cases.
Actual savings vary depending on architecture, region, volume and utilisation level.

AMINE RAITI · Infrastructure Architect & SRE · Paris

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

AMAZON WEB SERVICES — COMPUTE · NETWORK · EKS · VPN

Indicative estimates based on common use cases · Actual savings vary by architecture and volume

◆ BLOCK 1 · NAT GATEWAY — The most overestimated line item
HIGH IMPACT

Real cost: $0.045/h × 8,760h = ~$390/yr + $0.045/GB processed → up to $7,500/yr for a standard multi-AZ setup

- ✅ ALT 1 · NAT Instance (t3.nano/t3.small) → *same function · cost ~$130–220/yr · reduction 70–85% · trade-off: manual HA management*
- ✅ ALT 2 · VPC Endpoints (S3 · DynamoDB · SQS · SSM) → *transfers to internal AWS services are free · zero refactoring if already on AWS SDK*
- ✅ ALT 3 · PrivateLink for managed services → *avoids NAT Gateway transit · billed per connection only*
- 🔧 *When to keep NAT Gateway: critical workloads requiring managed high availability without human intervention*

◆ BLOCK 2 · EKS — Elastic Kubernetes Service
AVERAGE IMPACT

Real cost: $0.10/h per cluster = $876/yr per cluster · independent of node count · 5 clusters = $4,380/yr for control plane alone

- ✅ ALT 1 · Cluster consolidation: *5 env. clusters → 2 (prod + non-prod) with namespaces + Network Policies → direct control plane savings*
- ✅ ALT 2 · Karpenter vs Cluster Autoscaler → *provision at exact granularity required by pods · reduces over-provisioning 20–40%*
- ✅ ALT 3 · Spot instances for non-critical nodes (CI/CD · batch · dev) → *60–90% reduction on node cost*
- ✅ ALT 4 · k3s or k0s *self-managed on EC2 → eliminates control plane fees · for non-critical and edge environments*

◆ BLOCK 3 · EC2 RESERVED INSTANCES & SAVINGS PLANS
HIGH IMPACT

Trap: 1–3 year non-cancellable commitment · frequent over-commitment during growth peaks → fixed cost persists even if workloads disappear

- ✅ ALT 1 · Compute Savings Plans vs EC2 Instance SP → *more flexibility on instance type and region · preferred for evolving architectures*
- ✅ ALT 2 · Spot Instances for interruptible workloads → *60–90% reduction · ideal: ML training · batch · CI/CD · rendering*
- ✅ ALT 3 · Graviton (ARM) vs x86 → *same performance · 20% cheaper on equivalent instances · available on EC2 · Lambda · RDS · ElastiCache*
- 🔧 AWS Compute Optimizer → *rightsizing recommendations based on actual 14-day usage · free*

◆ BLOCK 4 · VPN & CONNECTIVITY
AVERAGE IMPACT

Direct Connect: billed by reserved capacity (1G ports = ~$1,500/mo) · often over-sized for light hybrid connections

- ✅ VPN Site-to-Site: $50–100/mo · *sufficient for <500 Mbps · 80–90% reduction vs Direct Connect*
- ✅ Direct VPC Peering (2–5 VPCs) → *free at the service level · only inter-region egress is billed · avoids Transit Gateway fees ($0.05/GB)*

◆ BLOCK 5 · ECS FARGATE
AVERAGE IMPACT

Fargate bills per vCPU + allocated memory = 3–4× more expensive than EC2 for stable, predictable workloads

- ✅ ECS on EC2 Spot *for predictable workloads → 50–70% reduction*
- ✅ Fargate Spot *for interruptible workloads → 50–70% reduction vs Fargate on-demand*
- ✅ Migrate to Kubernetes (EKS/k3s) *for organisations with K8s expertise → better cost control*

Amine RAITI · CC BY-NC-SA 4.0

AMAZON WEB SERVICES — S3 · RDS · LAMBDA · CLOUDWATCH · CLEANUP

Indicative estimates based on common use cases

◆ BLOCK 6 · S3 & DATA TRANSFER — HIGH IMPACT

Internet egress: $0.09/GB · Standard storage: $0.023/GB · storage tiers often misused

- ✅ S3 Intelligent-Tiering → auto-migration to IA/Glacier based on access · 40–70% reduction on cold data
- ✅ S3 Select → transfer only filtered data, not the entire object
- ✅ CloudFront in front of S3 → reduces direct requests and Internet egress
- ✅ MinIO S3-compatible on-premise for non-critical data → zero egress

◆ BLOCK 7 · RDS & DATABASES — HIGH IMPACT

RDS Multi-AZ db.r5.large: ~$440/mo · often over-sized by 2–3× vs actual usage

- ✅ Aurora Serverless v2 → billed per ACU · suited for variable workloads · scale-to-zero possible
- ✅ RDS Proxy → reduces connections · optimises utilisation · improves resilience
- ✅ PostgreSQL/MySQL on EC2 Spot + PgBouncer → 40–60% reduction · trade-off: manual maintenance
- ✅ Rightsizing via Performance Insights → identify over-provisioned instances

◆ BLOCK 8 · LAMBDA & SERVERLESS — AVERAGE IMPACT

Cold starts + duration billing + Provisioned Concurrency = frequent sources of unanticipated overspend

- ✅ Lambda Power Tuning → find the memory/cost optimum per function · open-source AWS tool
- ✅ Graviton2 for Lambda → 20% cheaper · 19% faster on average
- ✅ Continuous workloads → EC2 or Fargate often cheaper than Lambda beyond 15min/invocation
- ✅ SnapStart (Java) → eliminates cold starts on JVM functions

◆ BLOCK 9 · CLOUDWATCH · LOGS · OBSERVABILITY — AVERAGE IMPACT

Log ingestion: $0.50/GB · storage: $0.03/GB · can represent 15–20% of the total bill without monitoring

- ✅ OpenTelemetry + Grafana + Loki self-managed → 70–90% reduction on observability
- ✅ Subscription filters before CloudWatch ingestion → log only what is useful
- ✅ S3 + Athena for long-term log archiving vs CloudWatch Logs Insights
- ✅ Metric Streams to Grafana Cloud → cheaper than CloudWatch custom metrics

◆ BLOCK 10 · CLEANUP — AMIS · EBS · SNAPSHOTS · ELASTIC IPS · ECR · STOPPED INSTANCES — IMMEDIATE SAVINGS

- 🔵 Stale AMIs: each AMI stores underlying EBS snapshots ($0.05/GB/mo) · a CI/CD pipeline without cleanup accumulates hundreds of AMIs → monthly audit + retention policy (e.g. keep last 5)
- 🔵 Detached EBS Volumes: billed at $0.08–0.10/GB/mo even without an attached instance · identify via: aws ec2 describe-volumes --filters Name=status,Values=available
- 🔵 Orphaned EBS Snapshots: snapshots whose source volume has been deleted · no operational value · often hundreds of GB accumulated
- 🔵 Unassigned Elastic IPs: $0.005/h (~$44/yr each) · an EIP not attached to a running instance is billed
- 🔵 Untagged / old ECR images: no pull cost but storage at $0.10/GB/mo · ECR lifecycle policy for automatic cleanup
- 🔵 Stopped EC2 instances: attached EBS volumes continue to be billed · associated Elastic IPs too
- 🔧 Tools: AWS Cost Explorer (Idle Resources) · AWS Trusted Advisor · aws-nuke (full non-prod cleanup) · CloudCustodian · Infracost

Amine RAITI · CC BY-NC-SA 4.0

AMAZON WEB SERVICES — STEP-BY-STEP MIGRATION STRATEGY

Indicative estimates · Effort: L=Low (days) · M=Medium (weeks) · H=High (months)

◆ PROGRESSIVE MIGRATION PRINCIPLE

Do not migrate all at once. Identify high-egress or high contractual cost services first. Abstract SDK dependencies before physically migrating. Each service migrated to an open standard is an additional negotiation lever against AWS.

◆ BLOCK 11 · AWS MIGRATION STRATEGY → ALTERNATIVES

AWS SERVICE	ALTERNATIVE	METHOD	EFFORT	SAVING
S3	MinIO · Backblaze B2 · Cloudflare R2	Standard S3 SDK → zero refactoring	L	**Egress →0**
Lambda	OpenFaaS · Knative · Fission	Same Python/Node/Go runtime · adapt handler	M	**60–80%**
RDS PostgreSQL	Aiven · Neon · Render · Supabase	pg_dump / pg_restore · zero application refactoring	L	**30–50%**
EKS	k3s · k0s · EKS anywhere	Velero backup/restore · identical manifests	M	**Free Control plane**
CloudWatch	OpenTelemetry + Grafana + Loki	OTel agent → replaces CloudWatch agent	M	**70–90%**
Bedrock / SageMaker	vLLM + Llama 3 · Hugging Face self-hosted	OpenAI-compatible API → change endpoint URL only	M	**60–90%**
MSK (Kafka)	Apache Kafka self-managed · Redpanda	Same Kafka clients → zero refactoring	M	**40–60%**
ElastiCache Redis	Redis self-managed · KeyDB · Dragonfly	Same Redis protocol → zero refactoring	L	**40–60%**
NAT Gateway	NAT Instance t3.nano + VPC Endpoints	Replace in route tables · iptables	L	**70–85%**

Indicative estimates · L=days · M=weeks · H=months · discounts vs. AWS managed on-demand service

◆ RECOMMENDED MIGRATION ORDER

1. Observability (CloudWatch → OTel stack) — immediate bill impact
2. Object storage (S3 → MinIO/Backblaze B2) — zero refactoring
3. Relational databases (RDS → Aiven/Supabase) — pg_dump/restore
4. Serverless (Lambda → OpenFaaS) — adapt the handler
5. Kubernetes (EKS → k3s) — Velero for workloads
6. AI/ML (Bedrock → vLLM) — adapt the API endpoint

◆ AWS MIGRATION TOOLS

- ✅ *Velero → K8s workload backup/restore cross-cluster*
- ✅ *AWS DataSync → S3 data transfer to external destination*
- ✅ *OpenTofu (Terraform BSL fork) → portable multi-cloud IaC*
- ✅ *Infracost → cost estimates before Terraform deployment*
- ✅ *aws-nuke → full non-prod environment cleanup*

Amine RAITI · CC BY-NC-SA 4.0

GOOGLE CLOUD PLATFORM — BIGQUERY · GKE · COMPUTE · NETWORK

Indicative estimates based on common use cases

◆ BLOCK 1 · BIGQUERY — The most explosive line item

VERY HIGH IMPACT

Real on-demand cost: $5/TB scanned · a poorly written query on a 10TB table = $50 per run · a dashboard refreshing hourly = $1,200/mo

- ✅ Date/column partitioning → reduces scans by 80–95% on time-based queries · schema change only
- ✅ Clustering on frequently filtered columns → reduces scans on filter queries · stackable with partitioning
- ✅ Materialized views → pre-compute frequent aggregations · auto-updated · billed on delta only
- ✅ BI Engine → in-memory cache for dashboards · fixed monthly cost vs variable per-query cost · worthwhile from 3+ active dashboards
- ✅ Flat-rate slots vs on-demand → cost-effective from ~$2,200/mo in queries · 1-year or flexible monthly commitment
- ✅ Require partition filter → force queries to filter on partition · prevents accidental full scans
- ✅ Open-source alternatives: ClickHouse (OLAP columnar · 10–100× faster on aggregations) · DuckDB (local analytics · zero infra) · Apache Druid (real-time) · export BigQuery to Parquet via BQ Storage API → direct ClickHouse import with no transformation

◆ BLOCK 2 · GKE

AVERAGE IMPACT

Free control plane (1 Autopilot cluster) · nodes billed on requested resource

- ✅ Autopilot vs Standard → Autopilot bills pods · Standard bills full nodes · Autopilot better for variable workloads
- ✅ Spot VMs for non-critical nodes → 60–91% reduction
- ✅ Vertical Pod Autoscaler → auto-adjusts requests/limits · reduces over-provisioning
- ✅ KEDA → scale-to-zero for event-driven workloads

◆ BLOCK 3 · COMPUTE ENGINE & CUDS

HIGH IMPACT

CUD trap: 1–3 year commitment · no refund on termination (GCP s.8.8)

- ✅ Spot VMs → 60–91% reduction · ideal: ML training · batch · CI/CD
- ✅ Custom Machine Types → pay exactly the vCPU/RAM needed without rounding up
- ✅ Sole-tenant nodes for workloads requiring physical isolation · cheaper than over-specification
- 🔧 GCP Recommender → automatic rightsizing based on actual usage

◆ BLOCK 4 · VPC & NETWORK GCP

AVERAGE IMPACT

Cloud NAT: IP address + processed volume · Cloud Interconnect: reserved capacity (k$/mo)

- ✅ NAT on Spot instance → same approach as AWS · 70–85% reduction
- ✅ Cloud VPN (~$50/mo) vs Cloud Interconnect for hybrid connections <500 Mbps
- ✅ Direct VPC Peering vs Cloud Interconnect for simple architectures

◆ BLOCK 5 · CLOUD RUN & FUNCTIONS

AVERAGE IMPACT

Min-instances >0 = billed even without traffic · watch for idle costs

- ✅ Min instances = 0 for non-critical workloads → billed on usage only
- ✅ Cloud Run Spot → 60–91% reduction for batch processing
- ✅ Knative self-managed on GKE → removes Cloud Run dependency · same API
- ✅ CPU always allocated = false → CPU throttled outside requests = cheaper

Amine RAITI · CC BY-NC-SA 4.0

GOOGLE CLOUD — CLOUD STORAGE · CLOUD SQL · VERTEX AI · CLEANUP

Indicative estimates based on common use cases

◆ BLOCK 6 · CLOUD STORAGE & DATA TRANSFER — HIGH IMPACT

- *Internet egress: $0.08–0.12/GB by region · inter-region transfers: $0.01–0.08/GB*
- ✅ *Nearline / Coldline / Archive for cold data → 60–90% storage reduction*
- ✅ *Cloud CDN in front of GCS → reduces direct requests and egress*
- ✅ *Storage Transfer Service for migrations → reduced rate vs standard egress*
- ✅ *S3-compatible providers (Backblaze B2 · Cloudflare R2 · Wasabi) for non-critical data*

◆ BLOCK 7 · CLOUD SQL — HIGH IMPACT

- *HA instance db-n1-standard-4: ~$660/mo · storage and replicas billed separately*
- ✅ *AlloyDB for PostgreSQL-compatible workloads → better perf/cost on intensive loads*
- ✅ *Cloud SQL + pgBouncer → reduces connections · optimises resources*
- ✅ *PostgreSQL on GCE Spot + PgBouncer → 50–70% reduction · trade-off: maintenance*
- ✅ *Aiven for PostgreSQL · Neon · Supabase → managed · cheaper than Cloud SQL*

◆ BLOCK 8 · VERTEX AI & MANAGED AI — HIGH IMPACT

- *A100 80GB GPU: ~$3/h · V100: ~$2/h · training and inference cost highly variable by usage*
- ✅ *Spot VMs for training (interruptible by nature) → 60–91% reduction*
- ✅ *Model quantization + pruning → reduces model size and inference cost by 50–80%*
- ✅ *Hugging Face + vLLM self-hosted on Spot GPU → 60–90% reduction vs Vertex AI Prediction*
- ✅ *Ollama for local development → eliminates all inference costs in dev/test*

◆ BLOCK 9 · GCP OBSERVABILITY — AVERAGE IMPACT

- *Log ingestion: $0.50/GB · custom metrics: $0.10/metric/mo · can explode with many services*
- ✅ *OpenTelemetry + Prometheus + Grafana + Loki → full replacement · 70–90% reduction*
- ✅ *Log exclusion via Log Router → filter low-value system logs before ingestion*
- ✅ *GCP Managed Prometheus → cheaper than Cloud Monitoring custom metrics for K8s metrics*

◆ BLOCK 10 · CLEANUP — Persistent Disks · Snapshots · Images · Stopped Instances · Abandoned Buckets — IMMEDIATE SAVINGS

- 🔵 *Detached Persistent Disks: billed $0.04–0.17/GB/mo by type (Standard/SSD/Extreme) even without attached VM · identify via: gcloud compute disks list --filter="NOT users:*"*
- 🔵 *Orphaned snapshots: snapshots of deleted disks · no operational value · regular audit required · cost $0.026/GB/mo*
- 🔵 *Unused Compute Engine images: custom images unused for 90+ days → automatic archiving policy to Cloud Storage (Coldline)*
- 🔵 *Stopped instances: attached persistent disks continue to be billed · reserved unassigned IPs too ($0.010/h)*
- 🔵 *Abandoned Cloud Storage buckets: buckets created for completed projects · often forgotten · Object Lifecycle Management for automatic cleanup*
- 🔵 *Unmanaged Instance Groups: instances in unmanaged groups often forgotten after migrations · quarterly audit recommended*
- 🔧 *Tools: GCP Recommender (idle VMs · underused disks) · gcpdiag · Infracost · Active Assist · Cloud Asset Inventory for full resource audit*

Amine RAITI · CC BY-NC-SA 4.0

GOOGLE CLOUD — STEP-BY-STEP MIGRATION STRATEGY

Indicative estimates · Effort: L=Low (days) · M=Medium (weeks) · H=High (months)

◆ GCP-SPECIFIC WARNING — COMMITTED USE DISCOUNTS

GCP CUDs offer no refund on termination (s.8.8). Migration must be planned so commitments are consumed naturally during transition. Never terminate a CUD early — migration savings never offset the total loss of the committed amount.

◆ BLOCK 11 · GCP MIGRATION STRATEGY → ALTERNATIVES

GCP Service	Alternative	Method	Effort	Saving
BigQuery	ClickHouse · DuckDB	Export Parquet via BQ Storage API · direct import	M	**70–90%**
Cloud Storage (GCS)	MinIO · Scaleway Object	gsutil rsync → zero refactoring if S3 SDK abstracted	L	**Egress →0**
Cloud SQL PostgreSQL	Aiven · Clever Cloud	pg_dump / pg_restore · zero application refactoring	L	**30–50%**
GKE	k3s · k0s self-managed	Velero backup/restore · identical K8s standard manifests	M	**Free Control plane**
Cloud Functions	OpenFaaS · Knative	Same runtime · adapt the input handler	M	**60–80%**
Cloud Monitoring/Logging	OTel + Prometheus + Grafana + Loki	OTel Collector replaces GCP agent	M	**70–90%**
Vertex AI (inférence)	vLLM · Hugging Face	OpenAI-compatible API → change endpoint URL only	M	**60–90%**
Cloud Pub/Sub	Apache Kafka · Redpanda	Adapt publishers/ consumers · same semantics	M	**40–60%**
Memorystore Redis	Redis self-managed · KeyDB	Same Redis protocol → zero refactoring	L	**40–60%**

Indicative estimates · L=days · M=weeks · discounts vs. managed on-demand GCP service

◆ RECOMMENDED GCP MIGRATION ORDER

1. Observability (Cloud Monitoring → OTel stack) — immediate impact
2. Object storage (GCS → MinIO/Backblaze B2) — zero refactoring
3. BigQuery → ClickHouse — strong ROI if volume > 1TB/mo
4. Cloud SQL → Aiven/Neon — pg_dump/restore
5. GKE → k3s — Velero
6. Cloud Functions → Knative
7. Vertex AI → vLLM self-hosted

◆ GCP MIGRATION TOOLS

- ✅ *Velero → backup/restore K8s cross-cluster*
- ✅ *gsutil rsync → GCS transfer to S3-compatible destination*
- ✅ *BQ Storage API → export BigQuery to Parquet/Arrow*
- ✅ *Migrate for Anthos → GCE VM migration to containers*
- ✅ *OpenTofu → IaC portable multi-cloud*
- ✅ *gcpdiag → audit ressources orphelines GCP*

Amine RAITI · CC BY-NC-SA 4.0

MICROSOFT AZURE — AKS · NETWORK · BLOB · FUNCTIONS · RESERVATIONS

Indicative estimates based on common use cases

◆ BLOCK 1 · AZURE BLOB STORAGE & DATA TRANSFER — HIGH IMPACT

LRS: $0.018/GB · Internet egress: $0.07–0.09/GB · inter-region transfers: $0.02/GB

- ✅ *Archive tier → $0.001/GB storage · rehydration $0.02/GB · 94% reduction vs Hot*
- ✅ *Azure CDN in front of Blob → reduces direct requests and Internet egress*
- ✅ *Lifecycle Management → automatic archiving after X days without access*
- ✅ *MinIO S3-compatible on-premise for non-critical data → zero egress*

◆ BLOCK 2 · AKS — AZURE KUBERNETES SERVICE — AVERAGE IMPACT

Free control plane · Windows nodes 40% more expensive than Linux · billed on node uptime

- ✅ *Spot Node Pools → 60–90% reduction for non-critical workloads*
- ✅ *KEDA → scale-to-zero for event-driven workloads · évite les coûts idle*
- ✅ *Linux everywhere possible vs Windows → direct 40% saving*
- ✅ *Virtual Nodes (ACI) for bursting → no idle nodes · pay-per-pod*

◆ BLOCK 3 · VPN & EXPRESSROUTE — AVERAGE IMPACT

ExpressRoute: billed by reserved bandwidth · often over-sized for light hybrid connections

- ✅ *VPN Gateway Standard (~$140/mo) for hybrid connections <500 Mbps · 80–90% reduction vs ExpressRoute*
- ✅ *Direct VNet Peering vs Virtual WAN Hub → for simple architectures (2–5 VNets) · avoids Hub fees*
- ✅ *Azure Bastion vs VPN for admin access → fixed cost · more secure than VM-based VPN*

◆ BLOCK 4 · AZURE RESERVATIONS & SAVINGS PLANS — WARNING CONTRACT

⚠ Non-cancellable Savings Plans · early termination fee: 12% + $50,000/yr cap · Capacity Blocks non-refundable

- ✅ *Reserve only on the most stable and predictable resources · keep on-demand for variable resources*
- ✅ *Azure Reservations vs Savings Plans → Reservations more flexible on instance types for the same service*
- ✅ *Azure Hybrid Benefit → use existing Windows Server / SQL Server licences on Azure · 40–55% reduction*

◆ BLOCK 5 · AZURE FUNCTIONS & APP SERVICE — AVERAGE IMPACT

Consumption Plan · Premium Plan · Dedicated Plan → often sub-optimally chosen based on load pattern

- ✅ *Consumption Plan → pay-per-execution · ideal for sporadic workloads (<1M executions/mo)*
- ✅ *Premium Plan → pre-warmed instances · avoids cold starts · fixed billing · worthwhile for >8h/day activity*
- ✅ *Containers on AKS Spot for continuous workloads → more control · cheaper*
- ✅ *Knative on AKS → removes Azure Functions dependency · same scale-to-zero model*

◆ AZURE HYBRID BENEFIT — OFTEN UNDER-USED

If your organisation holds Windows Server or SQL Server licences with Software Assurance, Azure Hybrid Benefit allows reuse on Azure VMs or AKS. Reduction: 40–55% on Windows VM cost · up to 80% on Azure SQL database. Check systematically before any Windows workload deployment.

Amine RAITI · CC BY-NC-SA 4.0

MICROSOFT AZURE — SQL · COSMOS · OPENAI · SYNAPSE · MONITOR · CLEANUP

Indicative estimates based on common use cases

◆ BLOCK 6 · AZURE SQL & COSMOS DB — HIGH IMPACT

Azure SQL DTU: difficult to size · Cosmos DB RU: unpredictable cost without precise monitoring

- ✅ vCore vs DTU → *vCore more predictable · better cost/performance for stable workloads*
- ✅ Cosmos DB Autoscale throughput → *suited for variable workloads · avoids over-provisioning*
- ✅ Serverless Cosmos DB → *for sporadic workloads · billed on usage only*
- ✅ PostgreSQL + Citus (multi-tenant) → *complete open-source alternative to Azure SQL*
- ✅ FerretDB (MongoDB on PostgreSQL) → *Cosmos DB alternative for MongoDB APIs*

◆ BLOCK 7 · AZURE SYNAPSE & DATA FACTORY — AVERAGE IMPACT

Data Factory: billed per activity run · can explode on frequent pipelines without optimisation

- ✅ Pipeline consolidation → *group activities to reduce the number of billed runs*
- ✅ Self-hosted Integration Runtime → *run on existing machines · eliminates Azure IR fees*
- ✅ Serverless Synapse SQL Pool → *billed on queries only · no permanent nodes*
- ✅ Full open-source stack: Apache Airflow (orchestration) + dbt (transformation) + ClickHouse (storage)

◆ BLOCK 8 · AZURE OPENAI & COGNITIVE SERVICES — HIGH IMPACT

GPT-4: $30/M input tokens · GPT-4o: $5/M tokens · thousands of $/mo for intensive usage

- ✅ Caching frequent responses (Redis/Memcached) → 30–60% *reduction on repetitive queries*
- ✅ Model distillation → *fine-tune a smaller model on outputs of the large one* → 90% *inference cost reduction*
- ✅ Prompt optimisation → *reducing prompt size directly reduces the bill*
- ✅ Llama 3 / Mistral / Phi-3 on Spot GPU + vLLM → 60–90% *reduction vs Azure OpenAI*

◆ BLOCK 9 · AZURE MONITOR & LOG ANALYTICS — AVERAGE IMPACT

Log Analytics ingestion: $2.76/GB · retention beyond 31 days billed · Application Insights expensive at high volume

- ✅ OpenTelemetry + Grafana + Loki self-managed → 70–90% *reduction*
- ✅ Data Collection Rules (DCR) → *filter Windows/Linux events before ingestion*
- ✅ Basic Logs tier for low-value logs → *$0.60/GB vs $2.76/GB · 8-day retention*
- ✅ Log Analytics Archive tier → *$0.03/GB/mo after 90 days · on-demand queries*

◆ BLOCK 10 · CLEANUP — MANAGED DISKS · SNAPSHOTS · VM IMAGES · PUBLIC IPS · NSGS · APP REGISTRATIONS — IMMEDIATE SAVINGS

- 🔵 Unattached Managed Disks: *billed by type (Standard HDD $0.04/GB · Premium SSD $0.17/GB/mo) even without a VM · identify via: az disk list --query "[?diskState=='Unattached']"*
- 🔵 Orphaned snapshots: *snapshots of deleted disks · cost $0.04/GB/mo · monthly audit via Azure Cost Management*
- 🔵 Stale VM images: *generalised images unused for 90+ days → retention policy · archive to Azure Shared Image Gallery*
- 🔵 Unassigned Public IPs: *$0.0036/h (~$32/yr each) · an IP not attached to an active resource is billed*
- 🔵 Unused NSGs and routes: *no direct cost but increase attack surface and maintenance complexity · quarterly audit recommended*
- 🔵 Expired Entra App Registrations: *expired certificates and secrets · no direct cost but security risk and technical debt · Microsoft Entra Recommendations for audit*
- 🔵 Stopped (deallocated) VMs: *Managed disks continue to be billed · static public IPs too · delete or release attached resources*
- 🔧 Tools: *Azure Advisor (Cost tab) · Azure Cost Management + Budgets · AzureHound (IAM/resource audit) · Azure Resource Graph for cross-subscription audit queries*

Amine RAITI · CC BY-NC-SA 4.0

MICROSOFT AZURE — STEP-BY-STEP MIGRATION STRATEGY

Indicative estimates · Effort: L=Low (days) · M=Medium (weeks) · H=High (months)

◆ AZURE-SPECIFIC WARNING — SAVINGS PLANS & CAPACITY BLOCKS

Azure Savings Plans are non-cancellable (12% termination fee + $50,000/yr cap). Capacity Blocks are non-refundable. Plan migration so commitments are consumed naturally. Never terminate early without first calculating whether termination cost is lower than migration gain.

◆ BLOCK II · AZURE MIGRATION STRATEGY → ALTERNATIVES

Service Azure	Alternative	Méthode	Effort	Saving
Azure Blob Storage	MinIO · Backblaze B2 · Cloudflare R2	azcopy → rsync · S3-compatible SDK → zero refactoring	L	**Egress →0**
Azure Functions	Knative · OpenFaaS	Same Python/Node/Java runtime · adapt the handler	M	**60–80%**
Azure SQL	PostgreSQL managed (Aiven · Neon · Supabase)	pg_dump / pg_restore · zero application refactoring	L	**30–50%**
AKS	k3s · k0s · RKE2 self-managed	Velero backup/restore · identical K8s manifests	M	**Free Control plane**
Cosmos DB (MongoDB API)	MongoDB Community · FerretDB	Same MongoDB driver → zero refactoring if API compatible	L	**40–70%**
Azure Monitor / Log Analytics	OTel + Prometheus + Grafana + Loki	OTel Collector replaces Azure Monitor agent	M	**70–90%**
Azure OpenAI	Llama 3 · Mistral · Phi-3 + vLLM	OpenAI-compatible API → change endpoint URL only	M	**60–90%**
Azure Synapse / Data Factory	Apache Airflow + dbt + ClickHouse	Airflow DAG rewriting · dbt models	H	**50–70%**
Azure Cache for Redis	Redis self-managed · KeyDB · Dragonfly	Same Redis protocol → zero refactoring	L	**40–60%**

Indicative estimates · L=days · M=weeks · H=months · discounts vs. Azure managed on-demand service

◆ RECOMMENDED AZURE MIGRATION ORDER

1. Observability (Azure Monitor → OTel stack) — immediate impact
2. Blob Storage → MinIO/Backblaze B2 — zero refactoring
3. Azure SQL → managed PostgreSQL (Aiven/Neon) — pg_dump/restore
4. Azure Functions → Knative — adapt the handler
5. AKS → k3s — Velero for workloads
6. Cosmos DB → MongoDB Community — compatible API
7. Azure OpenAI → vLLM self-hosted — adapt endpoint

◆ AZURE MIGRATION TOOLS

- ✅ *Velero → backup/restore K8s cross-cluster*
- ✅ *azcopy → Blob transfer to S3-compatible*
- ✅ *Azure Migrate → VM assessment and migration*
- ✅ *AzureHound → IAM audit and orphaned resources*
- ✅ *OpenTofu → IaC portable multi-cloud*
- ✅ *Azure Resource Graph → cross-subscription audit queries*

Amine RAITI · CC BY-NC-SA 4.0

OPEN-SOURCE ALTERNATIVES BY CATEGORY — PORTABILITY & INDEPENDENCE

Indicative estimates · Portability: F=Full · P=Partial · Effort: L=Low · M=Medium · H=High

◆ TABLE 1 · OBSERVABILITY — Full replacement of CloudWatch / Cloud Monitoring / Azure Monitor — SAVING 70–90%

Composant	Role	Replaces	Portability
OpenTelemetry	Instrumentation & collect	CloudWatch Agent · GCP Agent · Azure Monitor Agent	T
Prometheus	Metrics	CloudWatch Metrics · Cloud Monitoring · Azure Metrics	T
Grafana	Dashboards & alerts	CloudWatch Dashboards · Google Cloud Dashboards · Azure Workbooks	T
Loki	Logs	CloudWatch Logs · Cloud Logging · Log Analytics	T
Tempo	Distributed traces	AWS X-Ray · Cloud Trace · Azure Application Insights	T

◆ TABLE 2 · OBJECT STORAGE — TOTAL PORTABILITY

Tool	Replaces	Effort
MinIO	S3 · GCS · Blob	L
Ceph / Rook	S3 · GCS · Blob (production)	M
Scaleway / OVHcloud S3	S3 · GCS · Blob (managé EU)	L

Standard S3 SDK (boto3 · minio-py) → zero application refactoring

◆ TABLE 3 · DATABASES — TOTAL PORTABILITY

Tool	Replaces	Effort
PostgreSQL + PgBouncer	RDS · Cloud SQL · Azure SQL	L
ClickHouse	BigQuery · Redshift · Synapse	M
MongoDB Community	DocumentDB · Cosmos DB	L
Redis / KeyDB / Dragonfly	ElastiCache · Memorystore · Cache for Redis	L

◆ TABLE 4 · ORCHESTRATION & K8s PORTABILITY

Tool	Role	Effort
k3s / k0s	Lightweight K8s · replaces EKS/GKE/AKS	M
Velero	K8s cross-cluster backup/restore	L
OpenTofu	Portable IaC (Terraform BSL fork)	L
Karpenter / KEDA	Intelligent Autoscaling K8s	M

◆ TABLE 5 · OPEN-SOURCE AI & ML — SAVING 60–90%

Tool	Replaces	Effort
vLLM	Bedrock · Vertex AI · Azure OpenAI	M
Ollama	Dev/test local · zéro coût cloud	L
Hugging Face + Ray Serve	SageMaker · Vertex AI · ML Studio	H
MLflow	SageMaker Experiments · Vertex Experiments	L

◆ TABLE 6 · SERVERLESS & EVENT-DRIVEN

Tool	Replaces	Portability	Effort
OpenFaaS · Knative	Lambda · Cloud Functions · Azure Functions	**Total after migration**	M
Apache Kafka · Redpanda	MSK · Pub/Sub · Event Hub · Event Bridge	**Total (same protocol)**	M
Apache Airflow	Step Functions · Cloud Workflows · Logic Apps	**Total**	M

Amine RAITI · CC BY-NC-SA 4.0

SUMMARY BY HYPERSCALER — LINE ITEMS · ALTERNATIVES · REDUCTIONS

Indicative estimates based on common use cases · Économies réelles variables selon architecture et volume

◆ AWS — SUMMARY BY IMPACT ORDER

Item	Main alternative	Saving est.	Effort
NAT Gateway	NAT Instance + VPC Endpoints	**70–85%**	**L**
EC2 Savings Plans	Spot + Graviton + Compute Optimizer	**20–90%**	**M**
RDS	Aurora Serverless v2 · PostgreSQL EC2 Spot	**30–60%**	**M**
CloudWatch Logs	OTel + Grafana + Loki	**70–90%**	**M**
EKS (multi-clusters)	Consolidation + Karpenter + Spot nodes	**40–80%**	**M**
Cleaning (EBS · AMIs · EIPs)	aws-nuke · Trusted Advisor	**Immediate**	**L**

◆ GCP — SUMMARY BY IMPACT ORDER

Item	Main alternative	Saving est.	Effort
BigQuery (queries)	Partitioning + Clustering + ClickHouse	**80–95% (scans)**	**L–M**
Compute Engine CUDs	Spot VMs + Custom Machine Types	**60–91%**	**L**
Cloud SQL	AlloyDB · PostgreSQL on Spot + PgBouncer	**30–70%**	**M**
Cloud Monitoring/Logging	OTel + Prometheus + Grafana + Loki	**70–90%**	**M**
Vertex AI (inference)	vLLM + open-source models on Spot	**60–90%**	**M**
Cleaning (Disks · Snapshots · Images)	GCP Recommender · gcpdiag	**Immediate**	**L**

◆ AZURE — SUMMARY BY IMPACT ORDER

Item	Main alternative	Saving est.	Effort
Azure OpenAI	vLLM + Llama 3 / Mistral on Spot	**60–90%**	**M**
Log Analytics	OTel + Grafana + Loki + Basic Logs tier	**70–90%**	**M**
ExpressRoute	Standard VPN Gateway	**80–90%**	**L**
Azure SQL / Cosmos DB	PostgreSQL + Citus · FerretDB	**30–70%**	**M**
Azure Hybrid Benefit	Use existing Windows/SQL licenses	**40–55%**	**L**
Cleaning (Disks · IPs · VM Images)	Azure Advisor · AzureHound	**Immediate**	**L**

All estimates are indicative and based on common use cases. Actual savings vary depending on architecture, region, volume, and usage level.
This document is purely technical. It contains no legal references or political dimension.
AMINE RAITI *· Infrastructure Architect & SRE · Paris · Public document · CC BY-NC-SA 4.0*
AI Powered by Amine — AI is an amplifier of ideas and forms, not the source.

Amine RAITI · CC BY-NC-SA 4.0

TECHNICAL GUIDE · INFRASTRUCTURE MIGRATION · JUNE 2026

VMWARE TO PROXMOX VE MIGRATION

Stretched Cluster D1 · D2 · Witness D3

Generic study · Full stack

From VMware/Broadcom lock-in analysis
to a 5-phase operational migration plan

6	3	5	7
VMware lock-in layers	Datacenters D1/D2/D3	Migration phases	Stack layers migrated

TABLE OF CONTENTS — 10 PAGES

◆ Page 2 · VMware/Broadcom lock-in anatomy
◆ Page 3 · Current VMware reference architecture
◆ Page 4 · Target Proxmox architecture
◆ Page 5 · Stretched cluster specific risks
◆ Page 6 · 5-phase migration plan
◆ Page 7 · Reference open-source stack
◆ Page 8 · Change management & skills
◆ Page 9 · Decision grid & approximate ROI
◆ Page 10 · Tools & resources

This document is a generic technical study. Configurations mentioned are illustrative.
It does not constitute individualised engineering advice. Public document · CC BY-NC-SA 4.0

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

VMware LOCK-IN — 6 LAYERS · BROADCOM CONTEXT

◆ BROADCOM POST-NOVEMBER 2023 CONTEXT — DOCUMENTED PRICING RUPTURE

Broadcom's acquisition of VMware in November 2023 immediately restructured licensing: perpetual licences discontinued · forced migration to annual subscriptions · documented price increases of ×3 to ×12 depending on configuration · Standard and Essentials editions eliminated in favour of imposed bundles. This is Layer 2 (pricing) of vendor lock-in applied to on-premise infrastructure — the same mechanism as cloud hyperscaler egress fees. The difference: the increase is immediate and without contractual alternative.

Layer	VMware/Broadcom mechanism	Intensity	Exit nature
1 · Contractual	Perpetual licences discontinued · non-cancellable subscriptions · support conditioned on active subscription · unilaterally modifiable Product Terms	●●●●	Legal · negotiation or migration
2 · Pricing	×3 to ×12 increase post-Broadcom · imposed bundles · double-run cost during migration · no alternative without full migration	●●●●	Financial · ROI <24 months in most cases
3 · Technical	VMDK format (convertible) · proprietary vSphere APIs · NSX-T network overlay · vSAN storage · Tanzu K8s · SPBM policies	●●●	Moderate technical · conversion tools available
4 · Skills	VCP/VCAP/VCDX certifications ubiquitous · attachment bias in ops teams · passive resistance from middle management	●●●	Training · reskilling plan
5 · Demand	vCenter as cognitive reference for infra teams · "we want vSphere" expressed by business units before any alternative evaluation	●●	Change management · functional reformulation
6 · Cognitive	Infra KPIs defined in vCenter/vRealize · Aria dashboards as performance benchmark · impossible to compare outside this benchmark	●●●	Irreversible without migration · rebuild KPIs

⚠ STRETCHED CLUSTER SPECIFICS — 3 ADDITIONAL COMPLEXITY DIMENSIONS

The D1/D2/D3 configuration adds three critical constraints absent from standard migrations: (1) vSAN inter-site synchronisation with tolerance for complete site loss — the D3 witness stores no data but arbitrates quorum; (2) NSX T0/T1 Gateway in active/standby mode across datacenters with North-South flow management between sites; (3) split-brain risk if the D1↔D2 link is cut — without D3 arbitration, both sites can diverge. These three dimensions impose a strict migration sequence: network before storage before compute.

Amine RAITI · CC BY-NC-SA 4.0

VMWARE STRETCHED CLUSTER ARCHITECTURE — D1 · D2 · D3

```
D1 (Primary Site)                 D2 (Secondary Site)               D3 (Witness)
├── ESXi Hosts (N)                ├── ESXi Hosts (N)                └── VMware Witness Appliance
├── vSAN datastore (preferred)    ├── vSAN datastore (secondary)        (vote only — zero data)
├── NSX Edge T0/T1 (active)       ├── NSX Edge T0/T1 (standby)
├── vCenter Server (active)       └── vCenter Server (passive HA)
└── DRS + HA enabled
         └────────── vSAN Stretched Cluster ──────────┘
              Synchronous replication D1 ↔ D2 · FTT=1 PFTT=1
              Max recommended latency: 5ms RTT D1↔D2 · 200ms D1/D2↔D3
```

◆ STORAGE LAYER — VSAN STRETCHED

vSAN Stretched Cluster with FTT=1 and PFTT=1: tolerates complete loss of one site (D1 or D2) without data loss. The Witness Appliance on D3 carries metadata components (witness components) to arbitrate quorum without storing VM data. SPBM storage policy: vSAN objects are distributed between D1 and D2 according to site-aware placement rules. On D1 loss, all VMs fail over to D2 via vSphere HA — the Witness D3 maintains quorum.

◆ NETWORK LAYER — NSX-T

NSX-T Tier-0 Gateway in active/standby inter-site mode: active T0 on D1 handles North-South routing · standby T0 on D2 takes over on D1 loss. Tier-1 Gateways per tenant/segment. GENEVE overlay on physical transport. Distributed Firewall (DFW) applied per VM vNIC — the firewall follows the VM regardless of physical location. VTEP on each ESXi host.

◆ ORCHESTRATION LAYER — VCENTER

vCenter Server in HA or Enhanced Linked Mode. DRS: workload balancing across hosts with site-aware affinity rules to avoid suboptimal inter-site placement. vSphere HA: automatic VM restart on surviving site on site loss. Fault Domain management: D1 = Preferred · D2 = Secondary.

◆ K8S LAYER — TANZU

VMware Tanzu Kubernetes Grid (TKG) or vSphere with Tanzu: K8s clusters supervised by vCenter · storage via vSAN CSI driver · networking via NSX-T NCP (NSX Container Plugin). Tanzu workloads co-located with VMs on the same vSAN/NSX fabric — Tanzu → K3s/K8s vanilla migration requires Velero for workload backup/restore.

◆ PROTECTION LAYER

Veeam Backup & Replication or VMware SRM for recovery plans. RPO/RTO defined by criticality tier: Tier 1 (RPO <15min · RTO <1h) · Tier 2 (RPO <1h · RTO <4h) · Tier 3 (RPO <24h · RTO <24h).

Amine RAITI · CC BY-NC-SA 4.0

PROXMOX STRETCHED ARCHITECTURE — CEPH · OVN · COROSYNC

```
D1 (Primary Site)                 D2 (Secondary Site)              D3 (Arbiter / Tiebreaker)
├── Proxmox VE Hosts (N)          ├── Proxmox VE Hosts (N)           └── Corosync QDevice (corosync-qnetd)
├── Ceph OSD (rack: dc1)          ├── Ceph OSD (rack: dc2)               Ceph MON tiebreaker
├── Ceph MON / MGR                ├── Ceph MON / MGR                     (vote only — zero OSD)
├── OVN Chassis (active)          ├── OVN Chassis (active)
└── Proxmox cluster node          └── Proxmox cluster node
         └──────── Ceph Stretched Cluster ───────────┘
              Synchronous replication D1 ↔ D2 · stretch_mode
              Max acceptable latency: 10ms RTT D1↔D2 · 30ms D1/D2↔D3
```

◆ STORAGE LAYER — CEPH STRETCHED

Ceph in stretched cluster mode with custom CRUSH map: bucket type `datacenter` containing D1 and D2 racks. CRUSH rule: `stretch_mode true · mon stretch_cluster · min_size=2 · size=4` (2 copies D1 + 2 copies D2). MON tiebreaker on D3 arbitrates Ceph quorum without OSD. RBD (RADOS Block Device) as VM storage backend — direct interface with Proxmox VE via Ceph pool.

◆ NETWORK LAYER — OVN + OVS

OVN (Open Virtual Network) with Open vSwitch: full functional replacement of NSX-T. Distributed OVN logical router: T0/T1 equivalent. GENEVE overlay (same protocol as NSX-T — same encapsulation). OVN ACLs: Distributed Firewall equivalent. BGP EVPN with FRRouting for inter-datacenter routing and D1↔D2 flow symmetry. OVS VTEP on each Proxmox host.

◆ ORCHESTRATION LAYER — PROXMOX CLUSTER

Proxmox VE cluster via Corosync + QDevice (corosync-qnetd) on D3. HA Groups per datacenter: `restricted=1` for Tier 1 VMs (stay on preferred site) · `restricted=0` for Tier 2/3 (auto-migrate to surviving site). Fence devices per site: IPMI/iDRAC for physical isolation on split. Live migration between D1 and D2 via KVM live migration (vMotion equivalent).

◆ K8S LAYER — K3S / K8S VANILLA

K3s or vanilla Kubernetes on Proxmox VMs · Longhorn or Ceph RBD CSI driver for persistent storage · Calico or Cilium for pod networking · Velero for workload backup/restore (migration from Tanzu). K3s recommended for small clusters (<20 nodes) · vanilla K8s for more complex environments.

◆ PROTECTION LAYER

Proxmox Backup Server (PBS): incremental VM backup with deduplication · configurable retention · client-side encryption. As complement or replacement: Veeam Community Edition (free up to 10 workloads) or existing Veeam licence during transition if contractually permitted.

vSAN function	Ceph equivalent	NSX-T function	OVN equivalent
Fault domain (site)	**CRUSH bucket datacenter**	T0/T1 Gateway	**OVN Logical Router**
Witness Appliance	**MON tiebreaker D3**	Distributed Firewall	**OVN ACLs**
PFTT=1 (site fault)	**stretch_mode min_size=2**	GENEVE overlay	**GENEVE (same protocol)**
SPBM policies	**Ceph pools + RBD**	NSX-T NCP (Tanzu)	**Calico/Cilium CNI**
vSAN Encryption	**Ceph at-rest encryption**	VTEP (per ESXi host)	**VTEP OVS (per PVE host)**

Amine RAITI · CC BY-NC-SA 4.0

CRITICAL RISKS — SPLIT-BRAIN · LATENCY · SEQUENCING

◆ RISK 1 — SPLIT-BRAIN: THE PRIMARY STRETCHED CLUSTER RISK

CRITICAL

If the D1↔D2 network link is interrupted, both sites may continue operating independently with diverging data. Without external arbitration (D3), both sites consider themselves the "live" site and continue accepting writes — producing irrecoverable data divergence.

VSAN STRETCHED
D3 Witness arbitrates via vSAN quorum. Site without quorum goes read-only. Automatic recovery on D1↔D2 reconnection.

CEPH STRETCHED
D3 MON tiebreaker arbitrates Ceph quorum. Site without quorum stops accepting writes (degraded PGs). Automatic recovery on reconnection.

COROSYNC/HA
D3 QDevice arbitrates Proxmox quorum. STONITH fencing per site. HA policy: fence minority site before VM relocation.

Critical prerequisite: the D1/D2↔D3 link must be independent of the D1↔D2 link. If D3 is on the same network path as D1↔D2, the tiebreaker cannot arbitrate a failure on that link.

◆ RISK 2 — INTER-SITE LATENCY: COMPARATIVE CONSTRAINTS

HIGH

Technology	Max recommended latency	Max acceptable latency	Impact if exceeded
vSAN Stretched D1↔D2	5ms RTT	10ms RTT	Degraded performance · SCSI timeouts
Ceph Stretched D1↔D2	10ms RTT	20ms RTT	OSD heartbeat warnings · I/O performance
Corosync/QDevice D↔D3	10ms RTT	30ms RTT	Unstable quorum · HA false positives
vMotion / Live migration	150ms RTT	250ms RTT	Migration duration increases
vSAN D1/D2 ↔ D3 Witness	200ms RTT	500ms RTT	Unstable vSAN quorum

Key point: Ceph stretched tolerates 10ms D1↔D2 latency vs 5ms for vSAN — making Proxmox/Ceph migration compatible with existing inter-site configurations that were at the vSAN tolerance limit.

◆ RISK 3 — MIGRATION SEQUENCING: MANDATORY ORDER

HIGH

A stretched cluster migration cannot be done layer by layer independently. The order is mandatory to avoid degrading fault tolerance during transition:

1. D3 infrastructure (QDevice + Ceph MON tiebreaker) → prerequisite for everything else
2. Network (OVN/OVS parallel to NSX-T) → dual-stack during transition
3. Storage (Ceph stretched parallel to vSAN) → VM data migration tier by tier
4. Compute (Proxmox VE parallel to ESXi) → VM migration after storage validated
5. K8s (K3s/K8s after compute) → Velero backup/restore from Tanzu
6. Protection (PBS replacing Veeam/SRM) → last, after full validation

Never decommission vSAN before Ceph stretched is validated in production. The storage double-run period is the most costly but least risky phase.

Amine RAITI · CC BY-NC-SA 4.0

5 PHASES · STRICT SEQUENCE · CONTROLLED DOUBLE-RUN

Durations are indicative estimates for a 20-host / 200-VM / 3-site configuration. Your context may vary significantly.

PHASE 0 · AUDIT & PREPARATION (estimate: 4–6 weeks)

CIO/CTO: full VM inventory · Tier 1/2/3 classification · double-run cost calculation · Broadcom contract audit · maintenance window identification · go/no-go decision per layer

SRE/Architect: D1↔D2↔D3 latency measurement · Ceph sizing (OSD/MON/MGR) · Proxmox VE sizing · OVN addressing plan · CRUSH rules definition · inter-site bandwidth validation

PHASE 1 · PARALLEL INFRASTRUCTURE (estimate: 2–4 weeks)

Proxmox VE installation on new hosts (or cohabitation if sufficiently dimensioned) · Corosync cluster D1+D2+QDevice D3 · Ceph MON tiebreaker D3 · Initial Ceph OSD D1 and D2 · Base OVN/OVS configuration · PBS installed and configured. **No VMs migrated.** Infrastructure validation phase only: simulated split-brain test · Corosync failover test · Ceph quorum test · latency validation.

PHASE 2 · TIER 3 MIGRATION (DEV/TEST) (estimate: 3–4 weeks)

VMDK → qcow2 conversion with `virt-v2v` or Proxmox import · cold migration of non-critical VMs · Ceph stretched performance validation under real load · live migration validation between D1 and D2 · Ceph failover validation on simulated site loss · team training on Proxmox interface. Network: first OVN segment in production parallel to NSX-T.

PHASE 3 · TIER 2 MIGRATION (NON-CRITICAL PRODUCTION) (estimate: 4–8 weeks)

Migration in batches of 10–20 VMs · Proxmox HA validation on simulated site loss · NSX-T to OVN flow migration segment by segment · OVN ACL validation as NSX DFW equivalent · Tier 2 K8s: Tanzu → K3s migration with Velero. Network: progressive NSX-T → OVN switchover with dual-stack maintained.

PHASE 4 · TIER 1 MIGRATION (CRITICAL PRODUCTION) (estimate: 4–8 weeks)

Planned migration with short maintenance windows (<30 min per batch) · final NSX-T → OVN switchover for critical flows · progressive vSAN deactivation after Ceph production validation · full disaster recovery test D1 → D2 · RPO/RTO validation per tier · documented rollback test. Tier 1 K8s: last migration after full Tier 2 validation.

PHASE 5 · VMWARE DECOMMISSION (estimate: 2–4 weeks)

Broadcom subscription termination · licence archiving · ESXi host removal · NSX-T manager decommission · vCenter deletion · configuration archiving · final Proxmox architecture documentation · actual vs estimated ROI audit · post-migration software governance audit.

◆ GOLDEN RULE FOR STRETCHED CLUSTER — NEVER DECOMMISSION BEFORE VALIDATION

vSAN must not be decommissioned before Ceph stretched has run in production across all migrated Tiers for at least 4 weeks without incident. NSX-T must not be decommissioned before OVN has managed all production flows for 2 weeks. The double-run cost is the price of safety.

Amine RAITI · CC BY-NC-SA 4.0

FULL MIGRATION TABLE — VMWARE → OPEN-SOURCE

Portability: F=Full (open standard) · P=Partial (adaptation required) · Effort: L=Low · M=Medium · H=High

Layer · VMware component	Open-source equivalent	Migration method	Effort	Portability
Hypervisor VMware ESXi + vSphere	Proxmox VE (KVM/QEMU) proxmox.com · AGPLv3	`virt-v2v` · OVA import · Proxmox import plugin · cold or live migration	M	**F · KVM standard**
Distributed storage vSAN Stretched	Ceph RBD Stretched ceph.io · LGPL	Export VMDK → qcow2 · RBD pool import · stretched CRUSH rules · D3 MON tiebreaker	H	**F · RADOS standard**
Network overlay NSX-T · GENEVE · DFW	OVN + Open vSwitch ovn.org · ovs.org · Apache 2.0	Segment-by-segment reconfiguration · OVN ACLs · BGP EVPN FRRouting · same GENEVE protocol	H	**F · open standards**
Cluster orchestration vCenter · DRS · HA	Proxmox Cluster Mgr + Corosync QDevice D3 tiebreaker	HA group reconfiguration · fence devices · migration policies · corosync-qnetd QDevice	M	**F · Corosync standard**
Kubernetes VMware Tanzu (TKG)	K3s / vanilla K8s k3s.io · kubernetes.io	Velero backup/restore · Ceph RBD CSI driver · Calico or Cilium CNI · cluster recreation	M	**F · CNCF standard**
Backup / DR Veeam · VMware SRM	Proxmox Backup Server (PBS) + Velero for K8s	PBS installation · job configuration · retention · encryption · full restore test	M	**F · open format**
Monitoring / Observability vRealize / Aria Operations	OTel + Prometheus + Grafana + Loki logs	OTel stack installation · Proxmox dashboards · Ceph dashboards · equivalent alerts	L	**F · CNCF standard**
IaC vSphere Terraform provider	OpenTofu + Proxmox provider opentofu.org · telmate/proxmox	Terraform → OpenTofu manifest rewrite · proxmox provider · same HCL logic	M	**F · HCL compatible**

◆ WHAT YOU GAIN BEYOND COST

Portable skills on open standards (KVM · Ceph · OVN · K8s) · Linux Foundation and CKAD/CKA certifications valued on the wider market · active community (Proxmox Forum · Ceph upstream · OVN GitHub) · enterprise support available (Proxmox GmbH ~€1,400/year/node · Red Hat Ceph Storage) · zero proprietary licence dependency to continue running infrastructure · EU Data Act art. 25 compatible: migration possible at any time without contractual obstacle.

Amine RAITI · CC BY-NC-SA 4.0

3 RESISTANCE PROFILES · TRAINING PLAN · CERTIFICATIONS

3 VMWARE-SPECIFIC RESISTANCE PROFILES — NATURE · LEVER · WHAT'S NOT SAID

① THE VCP/VCAP ENGINEER — VMWARE CERTIFIED FOR 3–8 YEARS

Resistance: "My VMware certifications are worthless if we migrate to Proxmox. I spent 3 years earning them."

Lever: KVM/Linux/Ceph/OVN skills are more portable and better compensated than proprietary VMware certifications. CKAD, CKA, Linux Foundation — larger market. Engage them as migration architect: they become the most valuable expert on the project.

What's not said: Their resistance is rational short-term. Over 3–5 years, open skills expose them to a much wider market than Broadcom certifications.

② THE NSX ARCHITECT — VMWARE NETWORK EXPERT

Resistance: "OVN is complicated, I know NSX-T inside out. The project will take twice as long if we change the network."

Lever: OVN uses the same GENEVE protocol as NSX-T. The distributed routing logic is identical. The learning delta is 6–8 weeks, not 6 months. Have them discover OVN alongside NSX-T during the dual-stack phase.

What's not said: Their NSX-T expertise transfers directly to OVN — the concepts are isomorphic. They're often the best OVN trainer on the team after 4 weeks of practice.

③ THE "VCENTER ADDICT" OPS — 5+ YEARS OF CONSOLE COMFORT

Resistance: "vCenter does everything in clicks. Proxmox is command line — we'll lose productivity."

Lever: Proxmox has a complete, intuitive web interface. Daily operations (create/clone/migrate a VM) are as fast in Proxmox GUI as in vCenter. Command line is a bonus, not a requirement for common ops.

What's not said: vCenter UX was designed to create this dependency. Naming it helps overcome it. A 30-minute live demo is usually enough to reduce the anxiety.

◆ TRAINING PLAN — 3 PHASES · SKILLS MAPPING

Current VMware skill	Open-source equivalent	Delta	Training duration
ESXi / vSphere admin	Proxmox VE + KVM/QEMU	**Medium**	3–4 weeks
vSAN admin	Ceph admin (OSD/MON/CRUSH)	**High**	6–8 weeks
NSX-T architect	OVN + OVS + FRRouting	**Medium**	4–6 weeks
Tanzu / vSphere K8s	K3s / K8s vanilla + Velero	**Medium**	3–5 weeks
Terraform vSphere provider	OpenTofu + Proxmox provider	**Low**	1–2 weeks

Amine RAITI · CC BY-NC-SA 4.0

AM I READY TO MIGRATE? · 12 QUESTIONS · SCORE /24

Decision support tool · Not exhaustive · ROI figures are approximate estimates for illustrative purposes only

Q.	Decision question · Yes = 2 pts · No = 0 pt	Layer
1	Has my Broadcom subscription increased by more than 50% since the VMware acquisition?	Pricing
2	Does my current contract include non-cancellable commitments beyond 12 months?	Contractual
3	Is my D1↔D2 latency below 10ms RTT (Ceph stretched compatible)?	Technical
4	Do I have an independent D3 site on a separate network path from D1↔D2 for the tiebreaker?	Technical
5	Does my team have at least 2 advanced Linux engineers capable of administering Ceph?	Skills
6	Do I have Tier 3 VMs (dev/test) without availability constraints to serve as a pilot?	Migration
7	Does my vSAN storage use complex SPBM policies that are difficult to reproduce in Ceph?	Technical
8	Are my NSX-T flows documented (T0/T1 topology, segments, DFW ACLs)?	Network
9	Using Tanzu: are my K8s workloads stateless or with documented persistent storage?	K8s
10	Do I have a budget for 6 months of double-run (VMware + Proxmox in parallel)?	Financial
11	Has my management approved a maintenance window for critical Tier 1 VMs?	Governance
12	Is my current disaster recovery plan documented and tested within the last 6 months?	DR

0–8 pts	9–14 pts	15–18 pts	19–24 pts
Insufficient prerequisites · In-depth audit before decision	Migration feasible · Start with Tier 3 · Resolve identified gaps	Migration recommended · 5-phase plan applicable	Priority migration · Fast ROI · Start Phase 0 immediately

◆ APPROXIMATE ROI — ILLUSTRATIVE EXAMPLE (20 HOSTS · 200 VMS · 3 SITES)

These figures are purely illustrative and approximate — your actual situation may differ significantly based on Broadcom negotiation, hardware sizing and internal migration costs.

Estimated VMware/Broadcom cost (illustrative): €300,000/year (post-Broadcom subscription ×3 on a previous base of €100,000/year).

Estimated Proxmox cost (illustrative): Optional Proxmox GmbH support ~€28,000/year (20 nodes × €1,400) · software licences: €0 (AGPLv3).

Estimated annual saving: ~€272,000/year after full migration.

Estimated migration cost: 6 months double-run (~€150,000 infrastructure) + training/consulting (~€50,000) = ~€200,000.

Estimated ROI: under 9 months after migration completion. These estimates are indicative — recalculate with your actual figures.

Amine RAITI · CC BY-NC-SA 4.0

MIGRATION TOOLS · DOCUMENTATION · COMMUNITIES

◆ VM CONVERSION TOOLS

VIRT-V2V — VMware VM conversion (VMDK) to KVM (qcow2) · supports ESXi, vCenter · automatic virtio driver handling · `libguestfs-tools`

QEMU-IMG — manual VMDK → qcow2 conversion · `qemu-img convert -f vmdk -O qcow2 vm.vmdk vm.qcow2`

PROXMOX IMPORT PLUGIN — direct import from ESXi datastore via Proxmox interface · available Proxmox VE 8.x

OVFTOOL — VMware export to OVA/OVF · used upstream of virt-v2v

◆ CEPH STRETCHED TOOLS

CEPHADM — Ceph deployment · recommended for Ceph 16+

CRUSHTOOL — CRUSH map compilation and validation

CEPH MON STRETCH CLUSTER — stretched mode activation command

RADOSGW-ADMIN — RGW object management if object storage required

◆ OVN NETWORK TOOLS

OVN-NBCTL / OVN-SBCTL — OVN North/South database management

OVS-VSCTL — Open vSwitch configuration

FRROUTING (FRR) — BGP EVPN for inter-site routing

PROXMOX SDN — Proxmox interface for OVN (simplifies configuration)

◆ K8S & BACKUP TOOLS

VELERO — K8s workload backup/restore · Tanzu → K3s/K8s vanilla migration · `vmware-tanzu/velero`

K3S — lightweight Kubernetes · ideal <20 nodes · `k3s.io`

CEPH CSI DRIVER — K8s persistent storage via RBD · `github.com/ceph/ceph-csi`

PBS (PROXMOX BACKUP SERVER) — incremental VM backup · deduplication · encryption · free

◆ REFERENCE DOCUMENTATION

PROXMOX VE ADMIN GUIDE — pve.proxmox.com/pve-docs/

CEPH DOCUMENTATION — docs.ceph.com · "Stretched Clusters" section

OVN ARCHITECTURE GUIDE — ovn.org/documentation

FRROUTING DOCS — docs.frrouting.org · BGP EVPN

RFC 7348 — VXLAN encapsulation

RFC 8365 — EVPN for datacenter interconnect

◆ COMMUNITIES & SUPPORT

PROXMOX COMMUNITY FORUM — forum.proxmox.com · very active · fast responses

CEPH MAILING LISTS — lists.ceph.io · ceph-users for operational questions

OVN GITHUB — github.com/ovn-org/ovn · issues and discussions

PROXMOX GMBH SUPPORT — ~€1,400/year/node · guaranteed SLA

RED HAT CEPH STORAGE — enterprise Ceph support

◆ KEY COMMANDS — QUICK START

```
# VMDK → qcow2 conversion
virt-v2v -i vmx vm.vmx -o local -os /var/lib/vz/images/100/
# Ceph stretched mode activation
ceph mon enable_stretch_mode <tiebreaker_mon> datacenter
# Corosync QDevice on D3
pvecm qdevice setup <qdevice-ip>
# Live VM migration D1 → D2
qm migrate <vmid> <target-node-d2> --online --with-local-disks 0
```

This guide is a generic technical study. Configurations are illustrative.
It does not constitute individualised engineering advice.

NEMO SUPRA LEGEM EST

Amine RAITI · Infrastructure Architect & SRE · Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

TECHNICAL TUTORIAL GUIDE · JUNE 2026

VMWARE TO PROXMOX VE MIGRATION

Complete Execution Tutorial · 13 Modules

Stretched Cluster D1 · D2 · Witness D3

From initial audit to continuous operations

Commands, configuration files and validation checklists

13	3	60+	4
Execution modules	Datacenters D1/D2/D3	Commands & configs	Failure scenarios

TECHNICAL FOLLOW-UP — COMPLEMENT TO THE ARCHITECTURE STUDY

This guide is the operational follow-up to the architecture study. Where the study formalises technical choices and the target trajectory, this tutorial provides the commands, configuration files and checklists required to execute every layer of the migration, from initial audit through to continuous operations.

All commands and configurations are generic and illustrative.
Adapt to your real environment before executing in production. Public document · CC BY-NC-SA 4.0

Amine RAITI · Infrastructure Architect & SRE

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine

TABLE OF CONTENTS

- ◆ Module 0 · Pre-Migration Audit and Inventory
- ◆ Module 1 · Sizing and Capacity Calculation
- ◆ Module 2 · Preparation and Pre-Flight
- ◆ Module 3 · Corosync Cluster and QDevice D3
- ◆ Module 4 · Ceph Stretched End-to-End
- ◆ Module 5 · SDN Network VXLAN/EVPN
- ◆ Module 6 · Security, Hardening and IAM
- ◆ Module 7 · HA, Fencing and Affinity
- ◆ Module 8 · v2v Migration Step by Step
- ◆ Module 9 · PBS Backup and Immutability
- ◆ Module 10 · Observability
- ◆ Module 11 · Failure Drills (Chaos Engineering)
- ◆ Module 12 · Continuous Operations and Acceptance

◆ HOW TO USE THIS GUIDE

Each module follows the same structure: objective, execution commands, complete configuration files, and a validation checklist at the end of the section. Modules are designed to be followed in order — each subsequent module assumes the previous module's validations are complete. Command blocks are generic and must be adapted to your organisation's naming conventions, IP address ranges and security conventions before execution.

This document is the operational follow-up to the reference architecture study. It does not replace support from an integrator or architect on a real high-stakes migration project — it provides the reproducible execution skeleton.

Amine RAITI · CC BY-NC-SA 4.0

PRE-MIGRATION AUDIT AND INVENTORY

Before any action on Proxmox, a thorough inventory of the existing VMware environment is the number one condition for success. Most documented migration incidents in the industry stem from an unidentified application dependency upstream — not from a technical flaw in the target platform.

0.1 · VM Inventory via PowerCLI

Full export script of the vCenter inventory to CSV, including virtual hardware configuration, SPBM storage policies and guest tools status.

```
Connect-VIServer -Server vcenter.local
Get-VM | Select-Object Name, NumCpu, MemoryGB,
  @{N="Datastore";E={($_ | Get-Datastore).Name -join ","}},
  @{N="SPBMPolicy";E={(Get-SpbmEntityConfiguration -VM $_).StoragePolicy.Name}},
  @{N="ToolsStatus";E={$_.ExtensionData.Guest.ToolsStatus}},
  @{N="OS";E={$_.ExtensionData.Guest.GuestFullName}},
```

0.2 · Criticality Classification (Tier 1/2/3)

1. **Tier 1 (Critical)** · RPO < 15 min, RTO < 1h. Transactional databases, ERP, payroll systems.
2. **Tier 2 (Important)** · RPO < 1h, RTO < 4h. Non-critical business applications, intranet.
3. **Tier 3 (Non-critical)** · RPO < 24h, RTO < 24h. Dev/test environments, sandboxes.

0.3 · NSX-T Network Dependency Mapping

Before replacing NSX-T with the Proxmox SDN, document the exact topology: logical segments, active DFW rules, T0/T1 gateways, anti-spoofing tables.

```
curl -k -u admin:password -X GET \
  "https://nsx-manager.local/policy/api/v1/infra/segments" \
  -H "Content-Type: application/json" > nsx_segments.json

curl -k -u admin:password -X GET \
  "https://nsx-manager.local/policy/api/v1/infra/domains/default/security-policies" \
```

0.4 · SPBM Storage Policy Audit

List the SPBM policies in use and their future mapping to Ceph CRUSH rules — this step determines the complexity of Module 4.

```
Get-SpbmStoragePolicy | Select-Object Name, AnyOfRuleSets, Description
```

0.5 · Windows Server Licensing Audit

For each Windows Server VM, check the licensing mode (per physical core, Datacenter, licence mobility) before migrating to KVM — a frequent and costly compliance gap if discovered after migration.

✓ **VALIDATION** — Complete CSV inventory generated and manually verified on 10% of VMs

✓ **VALIDATION** — Tier 1/2/3 classification validated with business teams

✓ **VALIDATION** — NSX-T export (segments + DFW) archived and documented

✓ **VALIDATION** — SPBM → CRUSH rules mapping table drafted

✓ **VALIDATION** — Windows Server licence compliance audit completed

Amine RAITI · CC BY-NC-SA 4.0

SIZING AND CAPACITY CALCULATION

1.1 · Ceph Sizing Formula — Usable Capacity

Usable capacity depends on the chosen replication factor. For a stretched architecture with size=4 (2 copies per site) and min_size=2:

```
Usable_capacity = Total_raw_capacity / replication_factor
Usable_capacity = Total_raw_capacity / 4
```

Example: 24 OSDs of 4 TB each = 96 TB raw → Real usable capacity = 96 / 4 = **24 TB exploitable**

1.2 · Required RAM per Ceph Node

Sizing rule: budget approximately 1 GB of RAM per TB of managed OSD, plus a 30% margin for recovery peaks.

```
Minimum_Ceph_RAM = (OSD_volume_per_node_in_TB × 1 GB) × 1.3
```

Example: a node with 4 OSDs of 4 TB = 16 TB per node → Minimum_Ceph_RAM = (16 × 1) × 1.3 = **20.8 GB dedicated to Ceph**

1.3 · Compute Capacity Calculation (vCPU/vRAM)

Method for calculating the target consolidation ratio while accounting for the HA margin (loss of one site).

```
Available_vCPU_per_site = (Physical_cores × Nodes_per_site) × oversubscription_ratio
Required_HA_margin = Total_VM_load / Active_sites_after_loss
```

With 2 active sites and tolerance for the loss of 1 site: each site alone must be able to absorb 100% of the total load.

1.4 · Reusable TCO Calculator

Line item	Formula	Variable to fill in
Hardware CAPEX	Nodes × Unit_cost	Vendor quote
Annual vendor support	Nodes × Support_cost	~€1,400/year/node (rough order)
Double-run transition	VMware_cost × Transition_months	Variable per existing contract
Net annual saving	VMware_cost - Proxmox_support	-

✓ **VALIDATION** — Ceph usable capacity calculation documented and validated per site

✓ **VALIDATION** — Ceph RAM calculated and compared against ordered hardware specs

✓ **VALIDATION** — HA margin paper-tested: a single site absorbs 100% of the load

✓ **VALIDATION** — TCO table filled in with the project's actual figures

Amine RAITI · CC BY-NC-SA 4.0

PREPARATION AND PRE-FLIGHT

2.1 · Proxmox VE Installation — Reproducible Deployment

Use the official ISO in graphical mode or via an answer file for automated, reproducible deployment across multiple nodes.

```
[global]
keyboard = "en-us"
country = "fr"
fqdn = "pve-d1-node1.lan.local"
timezone = "Europe/Paris"
root_password = "CHANGE_ME"

[network]
source = "from-dhcp"
```

2.2 · Network Configuration — Bonding and VLANs

Each node must have at least 2 network interfaces in active-active bonding (LACP) for Ceph back-end traffic, separated from the management network.

```
auto bond0
iface bond0 inet manual
    bond-slaves eno1 eno2
    bond-miimon 100
    bond-mode 802.3ad
    bond-xmit-hash-policy layer3+4

auto vmbr0
iface vmbr0 inet static
```

2.3 · Inter-Site Network Benchmark — iperf3

Validate actual latency and bandwidth between D1 and D2. Stretched Ceph tolerates up to 10ms RTT.

```
# On D2 (server)
iperf3 -s -p 5201

# On D1 (client) — 60 seconds, 4 parallel streams
iperf3 -c 10.20.1.11 -p 5201 -t 60 -P 4
```

✓ **VALIDATION** — Measured throughput ≥ 25 Gbps on the dedicated inter-site link

✓ **VALIDATION** — Average RTT latency < 10ms, stable jitter

2.4 · Raw Storage Benchmark — FIO

Raw performance test of NVMe disks before Ceph integration, to catch any underperforming disk before it becomes a cluster bottleneck.

```
fio --name=baseline_test --ioengine=libaio --rw=randwrite \
  --bs=4k --direct=1 --size=10G --numjobs=4 --iodepth=64 \
  --runtime=60 --group_reporting --filename=/dev/nvme0n1
```

✓ **VALIDATION** — Random 4k write IOPS consistent with manufacturer disk specs

✓ **VALIDATION** — Average write latency below 1ms on the disk alone

Amine RAITI · CC BY-NC-SA 4.0

PREPARATION AND PRE-FLIGHT

2.5 · Hardware Pre-Flight Checklist

✓ VALIDATION — HBA controllers in IT/Pass-through mode (no active hardware RAID)

✓ VALIDATION — Power Loss Protection (PLP) verified on every SSD destined for Ceph

✓ VALIDATION — Hardware watchdog active and tested (/dev/watchdog accessible)

✓ VALIDATION — BIOS set to performance mode, C-States disabled on compute nodes

OPS

Amine RAITI · CC BY-NC-SA 4.0

COROSYNC CLUSTER AND QDEVICE D3

3.1 · Forming the Proxmox Cluster

The first node initialises the cluster, subsequent nodes join it — creating the distributed configuration database (pmxcfs) shared across all nodes.

```
# On pve-d1-node1, initialisation
pvecm create cluster-stretched --link0 10.10.1.11

# On each subsequent node, joining
pvecm add 10.10.1.11 --link0 10.10.1.12
```

✓ **VALIDATION** — pvecm status shows the expected total number of nodes with quorum OK

3.2 · Deploying the QDevice on D3

The QDevice runs on an independent machine hosted at D3, fully separated from the D1-D2 network path.

```
# On the D3 machine
apt install corosync-qnetd -y

# On a Proxmox cluster node
apt install corosync-qdevice -y
pvecm qdevice setup 10.30.1.5 -f
```

✓ **VALIDATION** — pvecm status shows the QDevice as an active voting member

3.3 · Verifying Vote Distribution

With 3 nodes per site (D1 + D2) plus the QDevice, the total is 7 votes. Quorum is calculated automatically by Corosync on absolute majority.

```
pvecm status
# Expected output: Expected votes: 7, Total votes: 7, Quorum: 4
```

3.4 · Simulation Test — Loss of Inter-Site Link

Before going to production, simulate the D1-D2 link cut to validate the expected failover behaviour (the site with the QDevice obtains the majority).

```
# Network cut simulation on a test node
iptables -A INPUT -s 10.20.0.0/24 -j DROP
iptables -A OUTPUT -d 10.20.0.0/24 -j DROP
```

✓ **VALIDATION** — The priority site (with QDevice agreement) retains quorum

✓ **VALIDATION** — The isolated site triggers self-fencing (watchdog) without manual intervention

✓ **VALIDATION** — Restoring the link automatically restores the full cluster

3.5 · End-of-Module Checklist

✓ **VALIDATION** — 7-vote cluster operational and stable for more than 24h

✓ **VALIDATION** — Network cut test performed and documented with timestamp

✓ **VALIDATION** — Post-incident reconnection procedure validated for the ops team

Amine RAITI · CC BY-NC-SA 4.0

CEPH STRETCHED END-TO-END

4.1 · Initial Deployment via cephadm

cephadm is the recommended deployment tool for recent Ceph versions. It runs in a container and orchestrates installation across all nodes.

```
# Bootstrap the first monitor
cephadm bootstrap --mon-ip 10.10.1.11 --cluster-network 10.10.2.0/24

# Add hosts to the Ceph cluster
ceph orch host add pve-d1-node2 10.10.1.12
ceph orch host add pve-d1-node3 10.10.1.13
ceph orch host add pve-d2-node1 10.20.1.11
ceph orch host add pve-d2-node2 10.20.1.12
ceph orch host add pve-d2-node3 10.20.1.13
```

4.2 · Building the Per-Datacenter CRUSH Hierarchy

Create logical buckets of type datacenter to encapsulate each site's hosts — a prerequisite for enabling stretch mode.

```
ceph osd crush add-bucket D1 datacenter
ceph osd crush add-bucket D2 datacenter
ceph osd crush move D1 root=default
ceph osd crush move D2 root=default
ceph osd crush move pve-d1-node1 datacenter=D1
ceph osd crush move pve-d1-node2 datacenter=D1
ceph osd crush move pve-d1-node3 datacenter=D1
ceph osd crush move pve-d2-node1 datacenter=D2
ceph osd crush move pve-d2-node2 datacenter=D2
ceph osd crush move pve-d2-node3 datacenter=D2
```

✓ **VALIDATION** — ceph osd tree shows the D1/D2 hierarchy correctly populated

4.3 · Defining the Stretched CRUSH Rule

The rule enforces 2 copies on D1 and 2 copies on D2, guaranteeing full data survival in the event of a single-site loss.

```
ceph osd crush rule create-replicated stretched_rule default datacenter
ceph osd crush rule dump stretched_rule
```

4.4 · Activating Native Stretch Mode

Native stretch mode explicitly binds the D3 arbiter monitor and activates automatic failover behaviour.

```
ceph mon set election_strategy connectivity
ceph mon enable_stretch_mode pve-d3-mon stretched_rule datacenter
```

✓ **VALIDATION** — ceph -s shows the cluster in active "stretched" mode without warnings

Amine RAITI · CC BY-NC-SA 4.0

CEPH STRETCHED — RBD POOL AND FAILOVER TEST

4.5 · Creating the RBD Pool for Proxmox

```
# Create the VM storage pool
ceph osd pool create vm-storage 128 128 replicated stretched_rule
ceph osd pool set vm-storage size 4
ceph osd pool set vm-storage min_size 2
rbd pool init vm-storage

# Integration into Proxmox
pvesm add rbd ceph-stretched --pool vm-storage \
  --monhost "10.10.1.11 10.20.1.11 10.30.1.5"
```

4.6 · Real Failover Test — Simulated Loss of Site D1

Controlled shutdown of Ceph services on D1 to validate failover behaviour under realistic conditions.

```
systemctl stop ceph-osd.target
systemctl stop ceph-mon.target
```

✓ **VALIDATION** — ceph -s from D2 shows HEALTH_WARN with the pool degraded but writable

✓ **VALIDATION** — VMs on D2 continue operating without write interruption

✓ **VALIDATION** — On D1 service restart, recovery resynchronisation starts automatically

4.7 · End-of-Module Checklist

✓ **VALIDATION** — vm-storage pool operational with size=4, min_size=2 verified

✓ **VALIDATION** — D1 cut test performed successfully, D2 VMs uninterrupted

✓ **VALIDATION** — Post-incident resynchronisation time measured and documented

⚠ SRE ATTENTION POINT — INTER-SITE NETWORK LATENCY

Synchronous Ceph replication means a write operation incurs the round-trip inter-site latency (RTT). A 2ms RTT adds a fixed transaction cost of roughly 2ms. Dedicated fibre links and enterprise-class NVMe disks are essential to avoid degrading perceived application performance.

Amine RAITI · CC BY-NC-SA 4.0

SDN NETWORK VXLAN/EVPN

5.1 · Enabling the Proxmox SDN Module

```
apt install frr frr-pythontools -y
systemctl enable frr --now
```

5.2 · Configuring the EVPN Zone

An EVPN-type SDN zone encapsulates traffic in VXLAN with a BGP control plane for route distribution between D1, D2 and D3.

```
# /etc/pve/sdn/zones.cfg
evpn: zone-stretched
    controller bgp-controller
    vrf-vxlan 10000
```

5.3 · Defining Application VNets

Each VNet corresponds to a logical segment equivalent to an NSX-T segment, with an anycast gateway distributed across all hypervisors.

```
# /etc/pve/sdn/vnets.cfg
vnet: vnet-web
    zone zone-stretched
```

5.4 · Distributed Firewall (DFW Equivalent)

```
# /etc/pve/firewall/cluster.fw
[group sec-web-servers]
IN SSH(ACCEPT) -log info
IN HTTP(ACCEPT)
```

5.5 · Anti-Spoofing

```
pvesh set /nodes/pve-d1-node1/qemu/101/firewall/options --ipfilter 1 --macfilter 1
```

5.6 · Inter-Site Connectivity Test

```
vtysh -c "show bgp l2vpn evpn summary"
vtysh -c "show evpn vni"
```

✓ **VALIDATION** — EVPN routes learned on all three exit nodes (D1, D2, D3)

✓ **VALIDATION** — A VM on vnet-web at D1 successfully pings a VM on vnet-web at D2

✓ **VALIDATION** — Unauthorised inter-VNet traffic is blocked by the firewall

5.7 · End-of-Module Checklist

✓ **VALIDATION** — EVPN zone stable, BGP established between D1/D2/D3

✓ **VALIDATION** — Microsegmentation rules tested both positive and negative

✓ **VALIDATION** — Anti-spoofing active and verified on all critical VMs

Amine RAITI · CC BY-NC-SA 4.0

SECURITY, HARDENING AND IAM

6.1 · Hardening Cluster SSH Access

```
# /etc/ssh/sshd_config
PermitRootLogin prohibit-password
PasswordAuthentication no
PubkeyAuthentication yes
AllowUsers admin@10.0.0.0/8
```

6.2 · Enabling 2FA (TOTP) on the Web Interface

Proxmox VE natively supports TOTP for management interface access, independent of SSH authentication. Configuration via Datacenter > Permissions > Two Factor.

6.3 · LDAP/Active Directory Integration

Integrate the cluster with the existing corporate directory to centralise access management and traceability.

```
pveum realm add ad-company --type ad \
  --server1 ad.company.local \
  --domain company.local \
  --base-dn "DC=company,DC=local" \
  --bind-dn "CN=svc-proxmox,OU=Service,DC=company,DC=local"
```

6.4 · Managing Ceph and PBS Secrets

Ceph authentication keys (cephx) and PBS encryption keys must be stored outside the cluster itself, in a dedicated secrets vault.

```
ceph auth get client.admin -o /root/ceph-admin-key.backup
chmod 600 /root/ceph-admin-key.backup
```

6.5 · Permissions and Roles (RBAC) per Team

```
pveum role add VM-Operator -privs "VM.PowerMgmt,VM.Console,VM.Monitor"
pveum acl modify /vms/101 --roles VM-Operator --users ops-team@ad-company
```

6.6 · Log Auditing and Compliance

```
echo "*.* @@siem.company.local:514" >> /etc/rsyslog.conf
systemctl restart rsyslog
```

6.7 · End-of-Module Checklist

✓ **VALIDATION** — SSH password authentication disabled and keys deployed

✓ **VALIDATION** — 2FA active for all privileged administrator accounts

✓ **VALIDATION** — AD/LDAP integration functional, tested with a non-admin account

✓ **VALIDATION** — Ceph and PBS keys backed up outside the cluster in a dedicated vault

✓ **VALIDATION** — RBAC applied: no application account holds the full Administrator role

Amine RAITI · CC BY-NC-SA 4.0

HA, FENCING AND AFFINITY

7.1 · Configuring HA Groups

An HA group defines the set of nodes eligible to host a VM on failover, with an optional priority per node.

```
ha-manager groupadd group-tier1 \
```

7.2 · Adding a VM to HA Management

```
ha-manager add vm:101 --group group-tier1 --max_restart 3 --max_relocate 2
```

✓ **VALIDATION** — ha-manager status shows VM 101 as "started" with its group correctly assigned

7.3 · Anti-Affinity Rules

Prevent two redundant VMs (e.g. two domain controllers) from ending up on the same physical node after an automatic failover.

```
affinity_rules:
  - type: anti-affinity
```

7.4 · Hardware Watchdog and Self-Fencing

```
ls -la /dev/watchdog
```

✓ **VALIDATION** — The /dev/watchdog device is present and accessible by the Proxmox HA service

7.5 · Dynamic Load Balancing Script (DRS-like)

In the absence of a native dynamic scheduler equivalent to DRS, a scheduled script queries the Proxmox REST API to evaluate load and trigger live migrations.

```
import requests

def get_node_load(node):
```

7.6 · Real HA Failover Test

```
echo c > /proc/sysrq-trigger
```

✓ **VALIDATION** — The VM automatically restarts on an HA group node within 120 seconds

✓ **VALIDATION** — Anti-affinity rules are respected after the failover

7.7 · End-of-Module Checklist

✓ **VALIDATION** — All HA groups created and Tier 1/2 VMs correctly assigned

✓ **VALIDATION** — Watchdog tested and confirmed functional on every node

✓ **VALIDATION** — Node crash test performed with measured and documented RTO

Amine RAITI · CC BY-NC-SA 4.0

V2V MIGRATION STEP BY STEP

8.1 · Preparing the vCenter Connection in Proxmox

The native Proxmox import wizard connects directly to vCenter to list VMs available for migration, with no manual conversion step required.

```
pvesh create /nodes/pve-d1-node1/import/vmware \
  --vcenter vcenter.company.local --username svc-migration@vsphere.local
```

8.2 · Migrating a Linux VM — Worked Example

```
pvesh create /nodes/pve-d1-node1/import/vmware/import \
  --vm-id 201 --source-vm "app-linux-01" \
    --storage vm-storage --bridge vnet-web
```

✓ **VALIDATION** — The VM boots under Proxmox with disks accessible via RBD

✓ **VALIDATION** — The IP address and application services respond normally after boot

8.3 · Post-Migration Optimisation — VirtIO Drivers on Linux

```
lsmod | grep virtio
# Should show virtio_net, virtio_blk or virtio_scsi depending on disk config
```

8.4 · Migrating a Windows Server VM — Worked Example

The Windows procedure requires injecting VirtIO drivers before the first boot under Proxmox, or risk a blue screen at boot time.

```
# Run on the source Windows VM, before final cutover
pnputil -i -a "E:\viostor\2k22\amd64\viostor.inf"
pnputil -i -a "E:\NetKVM\2k22\amd64\netkvm.inf"
```

✓ **VALIDATION** — Successful boot with no INACCESSIBLE_BOOT_DEVICE error

✓ **VALIDATION** — Windows Device Manager reports no missing drivers

8.5 · Configuring vTPM for BitLocker

```
qm set 202 --tpmstate0 vm-storage:1,version=v2.0
```

✓ **VALIDATION** — The Windows VM detects the TPM 2.0 module in Device Manager

8.6 · Windows Licence Compliance Declaration

Document the underlying physical infrastructure change in the company's Software Asset Management (SAM) portal — a condition of Microsoft contractual compliance.

8.7 · End-of-Module Checklist

✓ **VALIDATION** — Pilot Linux VM migrated and validated in production over a full cycle

✓ **VALIDATION** — Pilot Windows VM migrated, VirtIO injected, vTPM active if applicable

✓ **VALIDATION** — SAM declaration completed for all migrated Windows Server VMs

Amine RAITI · CC BY-NC-SA 4.0

PBS BACKUP AND IMMUTABILITY

9.1 · Installing Proxmox Backup Server

PBS is installed standalone, ideally on infrastructure physically distinct from the production cluster to ensure isolation in case of compromise.

```
wget http://download.proxmox.com/debian/pbs/proxmox-release-bookworm.gpg \
```

9.2 · Creating the Backup Datastore

```
proxmox-backup-manager datastore create pbs-backup-pool /mnt/backup/pbs-store
```

9.3 · Connecting Proxmox VE to the PBS Datastore

```
pvesm add pbs pbs-storage --server pbs.company.local \
```

9.4 · GFS Retention Policy

```
# /etc/proxmox-backup/datastore.cfg
datastore: pbs-backup-pool
```

9.5 · Enabling WORM Immutability

Time-based locking on critical backups to prevent any deletion or modification, even by a compromised administrator account.

```
proxmox-backup-client protected-set <SNAPSHOT_ID> true \
```

9.6 · Replication to D3 (Pull Sync)

```
proxmox-backup-manager remote create pbs-d1 --host pbs.company.local --auth-id sync@pbs
```

9.7 · Full Restore Test

```
qmrestore <BACKUP_ID> 301 --storage vm-storage
```

✓ **VALIDATION** — Full restore completed within the target RTO defined for the VM's Tier

✓ **VALIDATION** — Restored VM boots and data integrity is confirmed

9.8 · End-of-Module Checklist

✓ **VALIDATION** — GFS policy applied and verified on the main datastore

✓ **VALIDATION** — WORM immutability enabled on Tier 1 backups

✓ **VALIDATION** — D3 replication functional, tested with a complete sync job

✓ **VALIDATION** — Full restore test documented with measured time

Amine RAITI · CC BY-NC-SA 4.0

OBSERVABILITY

10.1 · Installing the Collection Stack

VictoriaMetrics or Prometheus collect the metrics natively exposed by Proxmox VE and Ceph, with no additional agent on the hosts.

```
# /etc/pve/status.cfg
influxdb: metrics-server
    server metrics.company.local
    port 8089
    mtu 1450
```

10.2 · Ceph Exporter

```
ceph mgr module enable prometheus
ceph config set mgr mgr/prometheus/server_addr 0.0.0.0
ceph config set mgr mgr/prometheus/server_port 9283
```

✓ VALIDATION — curl http://localhost:9283/metrics returns Ceph metrics in Prometheus format

10.3 · Configuring Scrape Targets

```
# /etc/prometheus/prometheus.yml
scrape_configs:
  - job_name: 'ceph'
    static_configs:
      - targets: ['pve-d1-node1:9283', 'pve-d2-node1:9283']
  - job_name: 'proxmox'
    static_configs:
```

10.4 · Reference SLI/SLO Indicators

Indicator	Target threshold	Remediation action
Inter-site WAN latency	< 10ms (p99)	P1 alert, suspend live migrations
Ceph write latency	< 15ms (p99)	Check NVMe disks and back-end network
SSD wearout	< 85%	P3 alert, schedule preventive replacement

10.5 · Prometheus Alert Rules

```
groups:
  - name: ceph_alerts
    rules:
      - alert: CephOSDWearoutWarning
        expr: ceph_disk_wearout_ratio > 0.85
        for: 2h
        labels: {severity: warning}
```

10.6 · End-of-Module Checklist

✓ VALIDATION — Prometheus/VictoriaMetrics stack operational and collecting from all nodes

✓ VALIDATION — All SLI/SLO alert rules active and tested

✓ VALIDATION — Grafana dashboard validated with the ops team as a daily tool

Amine RAITI · CC BY-NC-SA 4.0

FAILURE DRILLS — LIGHTWEIGHT CHAOS ENGINEERING

The scenarios below should be executed in pre-production and then repeated in production during planned maintenance windows, at a minimum of one full cycle every 6 months.

11.1 · Scenario A — Total Loss of Datacenter D1

```
for node in pve-d1-node1 pve-d1-node2 pve-d1-node3; do
```

✓ **VALIDATION** — D2 retains quorum (3 votes + 1 QDevice = 4 votes out of 7)

✓ **VALIDATION** — Ceph switches to degraded single-site mode on D2 with no write loss

✓ **VALIDATION** — Tier 1 VMs from D1 automatically restart on D2 within 120 seconds

✓ **VALIDATION** — Automatic Ceph resynchronisation on return with no manual intervention

11.2 · Scenario B — Inter-Site Link Failure D1-D2

```
# Block traffic to the remote site only
```

✓ **VALIDATION** — The QDevice assigns its vote to the configured priority site (D1 obtains 4 votes)

✓ **VALIDATION** — Isolated D2 self-fences via watchdog without concurrent VM startup

✓ **VALIDATION** — No data corruption detected after the link is restored

11.3 · Scenario C — Single OSD Disk Failure

```
ceph osd out osd.5
```

✓ **VALIDATION** — Ceph automatically triggers rebalancing to the remaining OSDs

✓ **VALIDATION** — No application service interruption during rebalancing

✓ **VALIDATION** — Prometheus alert triggered in line with Module 10 rules

11.4 · Scenario D — Full Recovery After Major Disaster

Combination of the previous scenarios: complete loss of D1 AND a failed automatic failover, to validate the last-resort manual recovery procedure.

✓ **VALIDATION** — Manual procedure executable in under 4 hours (Tier 1 RTO)

✓ **VALIDATION** — Ops team capable of executing the procedure without the architect who designed the cluster

> ◆ **POST-DRILL REPORT**
>
> Each drill must produce a short report: actual failover duration, deviations from expected behaviour, identified corrective actions.

11.5 · End-of-Module Checklist

✓ **VALIDATION** — All 4 scenarios executed at least once in pre-production

✓ **VALIDATION** — Post-drill report written and corrective actions tracked

✓ **VALIDATION** — Recurring production drill calendar validated by operations management

Amine RAITI · CC BY-NC-SA 4.0

CONTINUOUS OPERATIONS AND ACCEPTANCE

12.1 · Updating the Proxmox VE Cluster

Updates are performed node by node, never simultaneously, to maintain quorum at all times during the operation.

```
pve6to7 --full
```

✓ **VALIDATION** — Each node rejoins the cluster with quorum intact before the next node is treated

12.2 · Updating Ceph

Strict order: monitors (MON) first, managers (MGR) next, then OSDs one at a time with a health check between each step.

```
ceph orch upgrade start --image quay.io/ceph/ceph:v18
```

✓ **VALIDATION** — ceph -s shows HEALTH_OK before moving to the next upgrade step

12.3 · Replacing a Failed Node

```
# Clean removal
ceph osd out osd.{id}
```

✓ **VALIDATION** — The new node joins the cluster with quorum, OSDs reintegrated and rebalancing complete

12.4 · Final Acceptance Checklist per Migrated VM

1. I/O performance comparable to or better than VMware (comparative FIO test)
2. All application services functional after a cold restart
3. PBS backup completed successfully at least once
4. HA failover test performed on this VM without error
5. Prometheus/Grafana monitoring active and alerts configured
6. Technical documentation updated (CMDB, network diagram)
7. Formal sign-off obtained from the business application owner

12.5 · Runbooks for Handover to Operations Teams

Standard structure: observed symptom · diagnostic checks · decision — automatic action vs manual intervention required · escalation procedure · reference architect contact.

12.6 · End-of-Module Checklist

✓ **VALIDATION** — Cluster upgrade procedure documented and tested in pre-production

✓ **VALIDATION** — Node replacement procedure tested at least once under real conditions

✓ **VALIDATION** — Acceptance grid applied to 100% of Tier 1 VMs before final cutover

✓ **VALIDATION** — Runbooks written for the 5 most probable incident scenarios

Amine RAITI · CC BY-NC-SA 4.0

TECHNICAL TUTORIAL · OPÉRATION DINDON · JUNE 2026

MARIADB
WITHOUT TREMBLING

Migration and Replication at Full Throttle

XtraBackup · SSH Streaming · Multi-TB with Zero Disk Saturation

◆ CONTEXT OF THIS TUTORIAL

Duplicating several terabytes of production data cannot be improvised. If you saturate your disk or lock your tables, you are no longer an engineer — you are a saboteur. This tutorial documents the complete method: direct SSH streaming with no local storage, data preparation, binlog and GTID replication, post-deployment verification. Clean and direct.

VOLUME	DOWNTIME	LOCAL STORAGE
Multi-TB	Zero	Zero

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · WHY XTRABACKUP — AND NOT THE OTHERS

CHOOSING THE RIGHT STRATEGY BEFORE TOUCHING ANYTHING

◆ MIGRATION STRATEGY COMPARISON

MYSQLDUMP / MARIADB-DUMP: simple, universal, human-readable. But it locks tables (or produces inconsistent data without locking), generates a SQL file to reimport line by line, and becomes impractical beyond a few dozen GB. Reserve for small databases or occasional migrations with no time constraint.

LVM SNAPSHOT: fast and consistent if the MySQL volume is on LVM. But requires prior LVM configuration, sufficient snapshot space, and a few-second maintenance window for the snapshot.

GALERA CLUSTER (SST/IST): ideal for already-running multi-master clusters. Irrelevant for a one-shot source → target migration.

XTRABACKUP / MARIABACKUP: physical hot backup, without locking InnoDB tables, direct streaming to the target, consistency guaranteed by transaction logs. This is the strategy for multi-TB production databases with zero downtime.

⚠️ XTRABACKUP VS MARIABACKUP — DO NOT CONFUSE

From MariaDB 10.2+, Percona XtraBackup is no longer fully compatible. Use **MARIABACKUP** — the official MariaDB fork. Commands are nearly identical but the binaries differ: `mariabackup` instead of `xtrabackup`. Check version: `mariadb --version`. This tutorial uses `xtrabackup` syntax for readability — adapt according to your installed version.

◆ WHY SSH STREAMING

SSH streaming pipes data directly from SOURCE to TARGET without ever writing to the SOURCE disk. Three critical advantages:

— **ZERO DISK SATURATION ON SOURCE:** a 3TB database does not require 3TB of free space on the source server.

— **REAL-TIME TRANSFER:** data arrives on the target as it is read, reducing total migration time.

— **SINGLE CHECKPOINT:** if the transfer fails, restart. No corrupted intermediate file to manage.

Prerequisites: passwordless SSH between the two servers (deployed public key) and sufficient network bandwidth between SOURCE and TARGET.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · PREPARATION — CHECKLIST BEFORE TOUCHING ANYTHING

ALIGNING THE STARS

✅ **PRE-MIGRATION CHECKLIST**

☐ Identical MariaDB versions on SOURCE and TARGET (`mariadb --version`)
☐ Identical XtraBackup/Mariabackup versions on both servers
☐ TARGET disk space ≥ SOURCE datafiles size + 20% margin (`du -sh /var/lib/mysql`)
☐ Passwordless SSH from SOURCE to TARGET and TARGET to SOURCE
☐ All critical tables are InnoDB (if MyISAM present, see warning below)
☐ Binlog enabled on SOURCE (`SHOW VARIABLES LIKE 'log_bin'`)
☐ GTID enabled or not — decide replication method before starting
☐ Firewall: MySQL port (3306 or custom) open between SOURCE and TARGET

⚠️ **MYISAM DANGER — READ BEFORE USING --NO-LOCK**

`--no-lock` is safe ONLY if all tables are InnoDB. A single MyISAM table produces silent inconsistency — the backup will appear to succeed but MyISAM data may be corrupted. Check: `SELECT TABLE_NAME, ENGINE FROM information_schema.TABLES WHERE TABLE_SCHEMA='your_db' AND ENGINE != 'InnoDB';` If non-empty result: convert to InnoDB or use `--lock-ddl-per-table` instead of `--no-lock`.

◆ CREATING USERS ON THE SOURCE

SQL — RUN ON SOURCE

```sql
-- User for XtraBackup (physical backup)
CREATE USER 'xtrabackup'@'localhost' IDENTIFIED BY 'SuperStrongPass';
GRANT RELOAD, LOCK TABLES, PROCESS, REPLICATION CLIENT
  ON *.* TO 'xtrabackup'@'localhost';

-- User for ongoing replication
CREATE USER 'repl_user'@'%' IDENTIFIED BY 'REPL_PASSWORD';
GRANT REPLICATION SLAVE ON *.* TO 'repl_user'@'%';
FLUSH PRIVILEGES;
```


Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · SSH STREAMING — ASPIRATION BY VACUUM

WE PIPE DATA DIRECTLY — NOTHING TOUCHES THE SOURCE DISK

◆ THE STREAMING COMMAND — RUN ON TARGET

BASH — RUN ON TARGET (IT PULLS DATA FROM SOURCE)

```
ssh root@source-db-server "bash -c '
  ulimit -n 1048576
  xtrabackup --backup \
    --user=xtrabackup \
    --password="SuperStrongPass" \
    --socket=/var/run/mysqld/mysql.sock \
    --no-lock \
    --parallel=4 \
    --stream=xbstream
'" | xbstream -x -C /backup/source-db-server
```

◆ OPTION-BY-OPTION BREAKDOWN

`ulimit -n 1048576`: raises the maximum number of simultaneously open file descriptors. On large databases with many tables, XtraBackup can hit the system default limit (1024) and fail silently. This line prevents it.

`--no-lock`: removes the global FLUSH TABLES WITH READ LOCK. Production continues uninterrupted. ONLY safe with 100% InnoDB (see Section 2).

`--parallel=4`: number of parallel read threads. Adjust based on available CPUs on SOURCE. Start at 4, increase if network bandwidth is not the bottleneck.

`--stream=xbstream`: XtraBackup proprietary streaming format, more efficient than tar for MySQL files.

`--socket`: path to the MariaDB Unix socket. Varies by distribution: `/var/run/mysqld/mysql.sock` (Debian/Ubuntu) or `/tmp/mysql.sock` (CentOS/RHEL). Verify: `SHOW VARIABLES LIKE 'socket';`

`xbstream -x -C /backup/source-db-server`: stream extraction on TARGET into the target directory. Ensure it exists and has sufficient space.

✅ MONITOR THE TRANSFER IN PROGRESS

```
# On TARGET — monitor progress
watch -n 5 "du -sh /backup/source-db-server"

# On SOURCE — monitor XtraBackup process
tail -f /var/log/mysql/xtrabackup.log
```


Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · PREPARE AND RESTORE — THE METAMORPHOSIS

MAKING DATA LIVE — THEN INJECTING IT

◆ STEP 1 — PREPARE: TRANSACTION LOG CONSOLIDATION

The physical backup captured in a live stream contains potentially inconsistent data pages — transactions open at capture time. `--prepare` applies the redo log and rolls back incomplete transactions, making the backup consistent and restorable. This step is mandatory. Without it, MariaDB will refuse to start on the restored files.

BASH — RUN ON TARGET

```bash
# Consolidation — makes the backup consistent
xtrabackup --prepare --target-dir=/backup/source-db-server

# Verify preparation completed without error
echo $? # must return 0
```

◆ STEP 2 — RESTORE: DEPLOYMENT ON TARGET

⚠️ **STOP MARIADB ON TARGET BEFORE RESTORE**

`systemctl stop mariadb` — The data directory must be empty. Never restore onto a running MariaDB instance.

BASH — RUN ON TARGET

```bash
# Empty the data directory (IRREVERSIBLE — double-check)
rm -rf /var/lib/mysql/*

# Restore files from the prepared backup
xtrabackup --copy-back   --target-dir=/backup/source-db-server   --datadir=/var/lib/mysql

# Fix permissions (MANDATORY)
chown -R mysql:mysql /var/lib/mysql

# Start MariaDB
systemctl start mariadb
systemctl status mariadb
```

◆ NASSIHA — DO NOT FORGET PERMISSIONS

XtraBackup restores files with the permissions of the user running the command (often root). Without `chown -R mysql:mysql`, MariaDB will refuse to start with a permissions error. This is the most frequent mistake at this stage.

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · REPLICATION — THE SACRED LINK

CONNECTING SOURCE AND TARGET — TWO METHODS

◆ RETRIEVING REPLICATION COORDINATES

XtraBackup created the file `xtrabackup_binlog_info` in the backup directory. It contains the exact binlog position at capture time.

BASH — READ COORDINATES

```bash
cat /backup/source-db-server/xtrabackup_binlog_info
# Example output:
mysql-bin.000123    456789    0-1-12345
# Column 1: binlog file
# Column 2: position (classic method)
# Column 3: GTID position (if GTID enabled)
```

◆ METHOD A — CLASSIC BINLOG (WITHOUT GTID)

SQL — RUN ON TARGET

```sql
STOP SLAVE;
CHANGE MASTER TO
  MASTER_HOST='source-db-server',
  MASTER_PORT=3306,
  MASTER_USER='repl_user',
  MASTER_PASSWORD='REPL_PASSWORD',
  MASTER_LOG_FILE='mysql-bin.000123',
  MASTER_LOG_POS=456789;
START SLAVE;
```

◆ METHOD B — GTID (RECOMMENDED MARIADB 10.2+)

More robust — survives restarts and failovers without manual position reconfiguration.

SQL — RUN ON TARGET

```sql
STOP SLAVE;
-- Inject GTID position from xtrabackup_binlog_info (3rd column)
SET GLOBAL gtid_slave_pos = '0-1-12345';
CHANGE MASTER TO
  MASTER_HOST='source-db-server',
  MASTER_PORT=3306,
  MASTER_USER='repl_user',
  MASTER_PASSWORD='REPL_PASSWORD',
  MASTER_USE_GTID=slave_pos;
START SLAVE;
```


Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · VERIFICATION — DO NOT LEAVE BEFORE CONFIRMING

SHOW SLAVE STATUS — FULL BREAKDOWN

◆ THE VERIFICATION COMMAND

SQL — RUN ON TARGET

```
SHOW SLAVE STATUS\G
```

`Slave_IO_Running: Yes` — the I/O thread is connected to SOURCE and reading the binlog. If `No`: network issue or incorrect credentials.

`Slave_SQL_Running: Yes` — the SQL thread is applying events on TARGET. If `No`: SQL error, check `Last_SQL_Error`.

`Seconds_Behind_Master: 0` — TARGET is synchronised. A non-zero value indicates replication lag — normal just after start, must converge to 0.

`Last_IO_Error / Last_SQL_Error` — if non-empty: read and address before continuing.

⚠️ COMMON ERRORS AND FIXES

Error 1236 — Could not find first log file: the binlog position in `xtrabackup_binlog_info` is stale (binlogs have been purged on SOURCE). Fix: purge_logs_before or redo the backup.

Error 1062 — Duplicate entry: a row already exists on TARGET. If non-critical: `SET GLOBAL slave_skip_errors = 1062;` and restart the slave. Investigate the root cause.

Error 2003 — Can't connect to MySQL server: check the firewall on SOURCE, bind-address in my.cnf (`bind-address = 0.0.0.0` or TARGET IP).

Slave_IO_Running: Connecting: incorrect credentials or repl_user not authorised from TARGET IP. Check `SHOW GRANTS FOR 'repl_user'@'%';`

RATIO

✅ POST-DEPLOYMENT MONITORING

```
# Monitor lag catchup in real time
watch -n 2 "mysql -e 'SHOW SLAVE STATUS\G' | grep -E 'Running|Behind|Error'"

# Verify consistency of a table between SOURCE and TARGET
# On SOURCE:
SELECT COUNT(*), MAX(id) FROM my_table;
# On TARGET (after Seconds_Behind_Master = 0):
SELECT COUNT(*), MAX(id) FROM my_table;
```

◆◆◆

If you saturate the disk or lock the tables, you are no longer an engineer — you are a saboteur. This strategy avoids both. Clean and direct.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART IX

LLM & ARTIFICIAL INTELLIGENCE

7 pieces

HUMAN

WORKING PAPER · INDEPENDENT CONTRIBUTION · JUNE 2026

FROM AMPLIFIER TO SUBSTITUTE

Towards a Theory of Individual Cognitive Lock-In Through Generative AI Coding Tools

STRUCTURED ABSTRACT

BACKGROUND

The mass adoption of generative AI coding tools (GitHub Copilot, ChatGPT, Cursor, Gemini Code) is reshaping software development practices at unprecedented speed. Existing literature focuses primarily on short-term productivity gains. Medium and long-term effects on developers' autonomous reasoning capacity remain underdocumented.

RESEARCH GAP

No formalised theoretical framework models the trajectory by which an initially amplifying use can shift towards structural dependency that atrophies autonomous code reasoning capacity.

CONTRIBUTION

This working paper proposes a four-phase model (Amplification · Drift · Structural Dependency · Individual Point of No Return) and a formal definition of individual cognitive lock-in as a phenomenon distinct from legitimate specialisation.

IMPLICATIONS

Proposals for organisations, educators and tool publishers. A longitudinal research protocol to empirically test the model.

Keywords: cognitive lock-in · generative AI · developer skills · code comprehension · technical debt · AI dependency · software engineering education · individual point of no return

Amine RAITI

Infrastructure Architect & SRE · Paris · Independent researcher

This document is an independent contribution to the scientific debate on the impact of generative AI tools on developer skills.
It is not part of any campaign or advocacy corpus. Tools mentioned as examples are cited as illustrations of the category.
It does not constitute legal advice. Public document · CC BY-NC-SA 4.0
AI Powered by Amine — AI is an amplifier of ideas and forms, not the source.

1. INTRODUCTION & CONTEXT

1.1 The adoption of generative AI coding tools

The adoption of generative AI coding tools has been one of the fastest transitions in the history of software engineering. GitHub Copilot, launched in general availability in June 2022, reached 1.3 million paying subscribers in 2023 and 55,000 enterprise customers. The Stack Overflow Developer Survey 2024 reveals that 76% of developers use or are considering using AI tools in their daily workflow. The category went from zero to ubiquitous in fewer than 36 months.

This mass adoption occurred without direct historical precedent. Previous transitions to new development paradigms — the shift to high-level languages in the 1960s, IDE adoption in the 1990s, framework proliferation in the 2000s — unfolded over decades, allowing the academic community time to document their effects on skills. Generative AI coding tools have not benefited from that delay.

1.2 Existing literature — focus on short-term productivity

Published studies converge on immediate productivity gains. Peng et al. (2023)[1] document a 55.8% increase in completion speed on the programming tasks defined in their experimental protocol. GitHub Research (2023)[2] measures a 56% increase in tasks completed per hour on standardised development tasks. These measurements are methodologically sound within their scope — they capture a real and significant short-term effect.

What these studies do not capture is equally significant: none measures the effect on autonomous reasoning capacity at 12, 18 or 24 months of intensive use. The temporal horizon of available studies generally stops at 4–8 weeks. That is sufficient to measure immediate productivity. It is insufficient to detect a progressive atrophy of reasoning capacity.

1.3 The theoretical gap

Three questions remain without formalised answers in existing literature: (a) How can one empirically distinguish amplifying use from substitutive use? (b) Is there an identifiable tipping point — a moment when the relationship to the tool changes qualitatively? (c) Is autonomous code reasoning capacity recoverable after a period of structural dependency, and at what cost?

1.4 Structure of the paper

Section 2 establishes the theoretical framework and distinguishes individual cognitive lock-in from legitimate specialisation. Section 3 presents the four-phase model. Section 4 examines the available evidence base and its limitations. Section 5 sets out implications for organisations, educators and tool publishers. Section 6 proposes a research agenda.

[1] Peng S. et al. (2023). *"The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." NBER Working Paper 31085.*
[2] GitHub (2023). *The Economic Impact of the AI-Powered Developer Lifecycle. GitHub Research Report.*

Amine RAITI · CC BY-NC-SA 4.0

2. THEORETICAL FRAMEWORK — SPECIALISATION VS COGNITIVE LOCK-IN

2.1 Legitimate specialisation — definition and reversibility criterion

A frontend developer who does not master the Linux kernel is specialised. This specialisation is legitimate: they have made a rational choice to concentrate their attention on a domain, and the reasoning capacity for lower layers is latent — available if the situation requires it, not exercised by choice. The operational criterion of specialisation is reversibility: if the developer decides to learn the Linux kernel, they can do so through standard learning effort. The learning capacity itself is not atrophied.

The theory of distributed cognition (Hutchins, 1995[3]) and the extended mind concept (Clark & Chalmers, 1998[4]) offer a framework for thinking about the integration of external tools into an individual's cognitive system. Clark & Chalmers argue that a notebook can be part of a person's cognitive system in the same way as their internal memory. This extension is cognitive but non-atrophying — removing the notebook reduces performance without degrading internal reasoning capacity.

2.2 Individual cognitive lock-in — formal definition

◆ FORMAL DEFINITION — INDIVIDUAL COGNITIVE LOCK-IN (ICL)

A state in which a developer can no longer reason autonomously and reliably about code they have themselves produced, without recourse to a generative AI tool, due to the progressive atrophy of their reasoning capacity through disuse.

Operational criteria: (a) inability to explain an algorithm committed within the last 6 months without AI-assisted regeneration; (b) inability to identify edge cases in a function produced without assisted generation; (c) dependency on AI-guided successive trials rather than causal reasoning for bug resolution.

2.3 Grounding in cognitive science literature

The theory of cognitive atrophy through disuse (Salthouse, 1991[5]) documents the degradation of unexercised cognitive capacities. Applied to code reasoning, it predicts that algorithmic reasoning capacity declines if it is not regularly solicited independently of the tool. The concept of cognitive offloading (Risko & Gilbert, 2016[6]) — the delegation of cognitive tasks to external artefacts — is relevant but incomplete: it does not model the irreversible degradation of the delegated capacity, only its temporary deactivation.

Individual cognitive lock-in is the case where cognitive offloading becomes irreversible — where the capacity delegated to the tool can no longer be recovered by simply deactivating the tool, but requires a relearning effort whose cost exceeds available motivation under continuous production pressure.

2.4 Structural analogy with vendor lock-in

The structure of ICL is analogous to that of cloud vendor lock-in (Opara-Martins et al., 2016[7]): an excellent tool in nominal cases · a dependency that builds progressively without a visible warning signal · a point of no return that makes the exit cost exceed available motivation. The analogy stops there: responsibility in ICL is tripartite (publisher · organisation · individual) whereas cloud vendor lock-in places responsibility primarily on the supplier side.

[3] Hutchins E. (1995). *Cognition in the Wild*. Cambridge: MIT Press.
[4] Clark A. & Chalmers D. (1998). "The Extended Mind." *Analysis* 58(1): 7–19.
[5] Salthouse T.A. (1991). *Theoretical Perspectives on Cognitive Aging*. Hillsdale: Lawrence Erlbaum Associates.
[6] Risko E.F. & Gilbert S.J. (2016). "Cognitive Offloading." *Trends in Cognitive Sciences* 20(9): 676–688.
[7] Opara-Martins J. et al. (2016). "Critical analysis of vendor lock-in and its impact on cloud computing migration." *JCSA* 5(4).

Amine RAITI · CC BY-NC-SA 4.0

3. FOUR-PHASE MODEL OF INDIVIDUAL COGNITIVE LOCK-IN

Calibrated on developers with intensive daily use of generative AI coding tools

PHASE 1 · AMPLIFICATION — 0 TO 6 MONTHS

Description: the developer uses AI to accelerate what they already know how to do. They understand the produced code, modify it with confidence, and can explain it during code review. Their judgment remains at the centre of the process. Their reasoning capacity is exercised.

Positive indicators: time to understand generated code is less than generation time · high post-generation modification rate · maintained explanatory capacity during code reviews · the tool reduces repetitive work without substituting judgment.

Defining characteristic: full reversibility — removing the tool does not degrade production capacity, only speed.

PHASE 2 · DRIFT — 6 TO 18 MONTHS

Description: natural cognitive resistance operates. If the tool produces something that passes automated tests, the effort of understanding encounters a growing motivational barrier. The developer begins committing code they only partially understand. The density of autonomous reasoning in their working day declines without a visible warning signal.

Drift indicators: time to understand generated code increases progressively · post-generation modification rate decreases · first difficulties explaining implementation choices during code reviews · first bug resolutions through successive trials rather than causal reasoning.

Triggering mechanism: short-term rational optimisation (accepting what works without understanding) produces a deferred cost invisible in standard productivity metrics.

PHASE 3 · STRUCTURAL DEPENDENCY — 18 MONTHS AND BEYOND

Description: the developer can no longer maintain their codebase without the tool. Documentable symptoms: inability to explain an algorithm committed several months earlier · inability to estimate the complexity of a modification without generating first · the AI context session becomes a necessary cognitive prosthesis · observable anxiety during sessions without access to the tool.

Organisational consequences: invisible technical debt (code works but nobody can explain or modify it with confidence) · fragility during production incidents (debugging without the tool equals paralysis) · aggravated bus factor (the developer depends on the tool, not only on their knowledge).

PHASE 4 · INDIVIDUAL POINT OF NO RETURN (PNR-I)

Formalisation: $PNR\text{-}i = \{t \mid C_recovery(t) > M_available(t)\}$ where C_recovery is the cognitive cost of recovering autonomous reasoning capacity (unlearning the dependency + relearning low-level reasoning) and M_available is the motivation available under continuous production pressure.

Dynamics: under continuous production pressure, M_available tends toward zero after 24–36 months of uninterrupted structural dependency. C_recovery increases with atrophy duration. The crossing of these two curves defines PNR-i. This model is proposed as a falsifiable hypothesis — its empirical validation is the object of the research protocol proposed in Section 5.

Amine RAITI · CC BY-NC-SA 4.0

4. AVAILABLE EVIDENCE BASE AND HONEST LIMITATIONS

4.1 What is documented

Several sources converge on signals consistent with the proposed model, without constituting direct proof of it.

◆ POSITIVE SIGNALS

Vaithilingam et al. (2022)[8] document that developers using generative AI tools frequently adopt code without fully understanding how it works. JetBrains Developer Ecosystem Survey 2024[9] signals growing concern among technical leads regarding the quality of AI-generated code. Thoughtworks Technology Radar Vol. 30 (2024)[10] introduces the notion of "AI-assisted development dependency" as an emerging risk.

⚠ WHAT IS MISSING

No rigorous published longitudinal study measures the degradation of autonomous reasoning capacity after 18–36 months of intensive substitutive use. The temporal horizon of available studies generally stops at 4–8 weeks. Existing observations are anecdotal or of short duration.

4.2 Assumed epistemic position

This working paper proposes a theoretical framework ahead of the available empirical base. This position is academically assumed and not concealed. It is consistent with the practice of working papers in cognitive science and innovation economics, where theorisation often precedes longitudinal data by necessity: phenomena with deferred effects cannot wait 24 months of data collection before being theorised.

The value of the proposed model is its falsifiability: the four phases, the associated indicators and PNR-i constitute testable hypotheses through a longitudinal protocol. If future data contradicts the model, the model must be revised. This is the condition of an honest working paper.

4.3 Proposed research protocol

◆ LONGITUDINAL DESIGN TO TEST THE MODEL

Cohort: 200 developers · varied experience levels · two groups (control with mandatory alternation policy · test with free use).

Measurements at baseline and T+6, T+12, T+18, T+24: (a) autonomous reasoning capacity on standardised algorithms; (b) debugging capacity without AI tool; (c) quality of explanations during simulated code reviews; (d) time to understand generated vs hand-written code; (e) accumulated technical debt measured by static analysis.

Testable hypotheses: H1 — The test group shows significant degradation of indicators (a)–(d) at T+18 vs baseline. H2 — Degradation is more pronounced in junior developers (<3 years experience). H3 — The mandatory alternation policy maintains the control group's indicators at a stable level over time.

[8] *Vaithilingam P. et al. (2022). "Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models." CHI Conference Extended Abstracts.*
[9] *JetBrains (2024). Developer Ecosystem Survey 2024. jetbrains.com.*
[10] *Thoughtworks (2024). Technology Radar Vol. 30. thoughtworks.com/radar.*

Amine RAITI · CC BY-NC-SA 4.0

5. PRACTICAL AND REGULATORY IMPLICATIONS

5.1 For organisations

◆ MEASURE 1 · CODE REVIEW POLICY

Obligation to explain during code reviews any AI-generated code. Not to prohibit — to require comprehension. A developer who cannot explain an algorithm they committed cannot validate it. This measure is the most directly effective because it forces the exercise of reasoning without excessive friction.

◆ MEASURE 2 · DELIBERATE ALTERNATION

One session per week without AI tools on critical parts of the codebase. Analogy: a musician who plays with an electronic metronome maintains practice sessions without the metronome to preserve their internal tempo capacity.

◆ VIGILANCE INDICATOR

If a developer cannot explain an algorithm they committed 6 months ago, the drift is active. Signal for a training intervention — not a sanction. Drift is a natural response, not a fault.

5.2 For educators and engineering schools

Algorithmic fundamentals must be maintained as a non-negotiable baseline, independent of available tools. Generative AI should be introduced as a second-level tool — after mastery of fundamentals — not as primary access to code. The pedagogical analogy is mathematics: one does not learn arithmetic with a calculator before understanding operations. The calculator amplifies an already present skill. It cannot create an absent one.

A pedagogical research direction: design evaluations that deliberately test reasoning capacity without AI tools, regularly and over the duration of training programmes. Not as a sanction of the tool, but as a measure of latent capacity.

5.3 For tool publishers

Design mechanisms that favour understanding rather than simple acceptance of generated code are both ethically justified and commercially rational in the long term: a developer who understands the code they generate with the tool is a more durable and loyal user than a developer who depends on it structurally without understanding.

◆ THREE DESIGN PROPOSALS

1. Mandatory explanation before acceptance: before accepting a generated code block, the tool requests a natural language explanation of what the code does. This forces active reading.
2. Learning mode: an option that reveals the step-by-step reasoning underlying the generation, allowing the developer to understand the logic, not just the result.
3. Dependency alert: if the acceptance-without-modification rate exceeds a threshold over a given period, the tool alerts the developer — not to penalise, but to signal a possible drift toward Phase 2.

5.4 Open regulatory question

Structural cognitive dependency on a proprietary tool raises an unresolved legal question: does it fall under employment law (skills required for the position) · competition law (entry barriers created by tool dependency) · or professional training law? No existing framework explicitly covers this case. It is an open legal research territory that goes beyond the scope of this working paper.

Amine RAITI · CC BY-NC-SA 4.0

6. CONCLUSION & FIVE PRIORITY QUESTIONS

6.1 Summary

This working paper proposes the first formalised theoretical framework for individual cognitive lock-in through generative AI coding tools. It distinguishes this phenomenon from legitimate specialisation through the criterion of reasoning capacity reversibility. It models four phases with operational indicators for each. It formalises an individual point of no return (PNR-i) analogous to the economic point of no return of cloud vendor lock-in.

The principal contribution is not empirical — the evidence base is explicitly limited — but theoretical: providing a framework that makes the phenomenon observable, measurable and therefore falsifiable. Without a framework, the early signals of drift remain invisible. With a framework, organisations and researchers can design indicators, measurement protocols and preventive interventions.

6.2 The author's position

◆ METHODOLOGICAL DECLARATION

This document is not an anti-AI argument. The author uses generative AI tools daily in their work as an architect and researcher. The position maintained deliberately is that of the amplifier: "*AI is an amplifier of ideas and forms, not the source.*" The ideas in this document are the author's. AI tools amplified the form. The question this document raises is not "should one use AI?" — the answer is yes. The question is: "what relationship does one maintain with this tool, and is that relationship chosen or undergone?"

6.3 Five priority research questions

Q.	Priority question
Q1	Is PNR-i empirically measurable and at what temporal horizon? The model predicts 24–36 months — is this confirmed by longitudinal data?
Q2	Is individual cognitive lock-in reversible with a targeted intervention, and at what cognitive and temporal cost?
Q3	Are there developer profiles more resistant to Phase 2 drift — and what are the predictors (prior experience · training type · code review practices)?
Q4	Is the mandatory alternation policy (AI-free sessions) sufficient to maintain autonomous reasoning capacity over the long term?
Q5	Is the phenomenon specific to programming or generalisable to other intellectually intensive professions using generative AI (writing · legal analysis · academic research)?

This document is an independent contribution. It does not constitute legal advice or professional counsel.
The author has no affiliation with the companies or institutions mentioned.
PER ASPERA AD ASTRA
Amine RAITI · Infrastructure Architect & SRE · Paris · Public document · CC BY-NC-SA 4.0
AI Powered by Amine — AI is an amplifier of ideas and forms, not the source.

REFERENCES

Academic literature

Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.

Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge: MIT Press.

Opara-Martins, Justice, Reza Sahandi, and Feng Tian. 2016. "Critical Analysis of Vendor Lock-In and Its Impact on Cloud Computing Migration: A Business Perspective." *Journal of Cloud Computing: Advances, Systems and Applications* 5 (4).

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." NBER Working Paper 31085. Cambridge: National Bureau of Economic Research.

Risko, Evan F., and Sam J. Gilbert. 2016. "Cognitive Offloading." *Trends in Cognitive Sciences* 20 (9): 676–688.

Salthouse, Timothy A. 1991. *Theoretical Perspectives on Cognitive Aging*. Hillsdale: Lawrence Erlbaum Associates.

Vaithilingam, Priyan, Tianyi Zhang, and Elena L. Glassman. 2022. "Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models." *CHI Conference on Human Factors in Computing Systems — Extended Abstracts*. New Orleans.

Industry and institutional reports

GitHub. 2023. *The Economic Impact of the AI-Powered Developer Lifecycle*. San Francisco: GitHub Research.

JetBrains. 2024. *Developer Ecosystem Survey* 2024. Prague: JetBrains. jetbrains.com/lp/devecosystem-2024.

Stack Overflow. 2024. *Developer Survey* 2024. New York: Stack Overflow. survey.stackoverflow.co/2024.

Thoughtworks. 2024. *Technology Radar Volume* 30. Chicago: Thoughtworks. thoughtworks.com/radar.

Note on tools cited as examples

GitHub Copilot (Microsoft), ChatGPT (OpenAI), Cursor (Anysphere) and Gemini Code (Google) are mentioned solely as illustrations of the "generative AI coding tools" category. Their mention constitutes neither a specific critique nor a recommendation. Other tools in the same category present the same structural characteristics.

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE REPLACEMENT THAT REVEALS

Why AI makes humans more precious

From the scarcity of lived experience to the trust crisis

HUMAN

◆ CONTEXT OF THE STUDY — IN DIALOGUE WITH ANOTHER VOICE

This study was born from a LinkedIn post by Sirine Houara — an engineer at the intersection of tech, marketing and human sciences — who formulated something this corpus had not yet explicitly stated: AI does not replace humans. It reveals their scarcity. The more generated content there is, the more authentic human experience becomes a rare and sought-after value. This thesis is the content-side mirror of what the corpus said about infrastructure. The mechanism is the same. The consequences differ.

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

I

SECTION I · THE CENTRAL PARADOX

AI MAKES HUMANS MORE VALUABLE — BY MAKING THEM SCARCE

The inversion is counter-intuitive — and that is precisely why it deserves to be stated clearly. We expected AI to replace humans. What it produces instead is less predictable: it reveals the scarcity of what it cannot reproduce. And in doing so, it increases the value of what it was supposed to devalue.

◆ THE ECONOMIC MECHANISM OF SCARCITY THROUGH ABUNDANCE

Basic economics: when something becomes abundant, its value falls. When something becomes scarce, its value rises. AI has made content production abundant — anyone can produce any text in seconds. Well-written, coherent, informative content is no longer a differentiating advantage. It is the norm. What remains scarce — and whose value is rising — is precisely what AI cannot produce: the irreproducible lived experience, the authentic contradiction, the error that reveals thought, the conviction carried by a real story.

◆ THE PARALLEL WITH THE INFRATIONAL LOOP

This mechanism is exactly the one the corpus documented in "The Infrational Loop" and "The Infrational Crisis" — but applied to content rather than technical skills. DevOps diluted bare-metal competences by making them vague and universal → real bare-metal competence became scarce → its value rose. AI dilutes content production by making it universal → authentic human content becomes scarce → its value rises. Same mechanism. Two different domains. The corpus documented infrastructure. Sirine documented content. The thesis is the same.

◆ NASSIHA — THIS IS NOT A CONSOLATION FOR HUMANS

This thesis is not an easy consolation of the type "don't worry, humans will always be useful". It is an economic mechanism analysis. The value of human singularity is rising — but only for those who genuinely have singularity to offer. Those who used their content production capacity as their main advantage without real singularity behind it — those people are indeed being replaced. The distinction matters.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · WHAT AI CANNOT GENERATE

THREE IRREDUCIBLES OF HUMAN SINGULARITY

AI can imitate a style. It can synthesise knowledge. It can generate coherent, informative, well-structured texts. What it cannot generate is what makes a human irreducible to their production function — the three components of singularity that Sirine identified in her post.

◆ LIVED EXPERIENCE — THE IRREPRODUCIBLE TRAJECTORY

AI can write about loss, resilience, failure, joy. It can do so with precision and even a certain elegance. What it cannot do is have lived through these experiences. The difference is not stylistic — it is ontological. When Sirine writes about her journey at the intersection of tech, marketing and human sciences, she mobilises an accumulation of real experiences, imperfect choices, encounters, errors and reconsiderations spanning years. This trajectory cannot be reproduced by a language model — because it was never data to train on. It is unique because it was lived, not because it is well written.

◆ CONTRADICTION — THE SIGN OF LIVING THOUGHT

An LLM optimises for coherence. It produces internally consistent texts without contradictions, without position reversals, without admissions of past error. This is precisely its weakness — because living thought is recognised by its capacity to contradict itself, to evolve, to acknowledge it was wrong. Contradiction is not a defect of human thought — it is its signature. An author who has never contradicted themselves has not thought much. An author who visibly evolves over years creates an intellectual trajectory that nobody can imitate without having genuinely lived through the same experiences.

◆ EMBODIMENT — PHYSICAL PRESENCE IN WHAT ONE PRODUCES

The dyslexic's spelling errors are more authentic than the model's perfect text. This is not a defence of carelessness — it is an observation about authenticity signals. Embodiment is the presence of a real person in what they produce — their hesitations, their recurring formulas, their emphases, their avoided subjects, their compulsive returns to the same themes. It is the corpus of a life that shows through in the corpus of a work. AI produces texts without a body, without history, without obsessions. It cannot embody itself because it has no body to embody.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE TRUST CRISIS — THE CENTRAL FILTER

IN A WORLD SATURATED WITH GENERATED CONTENT, TRUST COMES FIRST

Sirine's question — "In a world saturated with generated content, what is your trust based on today?" — is not a rhetorical question. It is the structuring question of the coming years. It points to a documentable problem: how do you distinguish the authentic from the generated when both are stylistically indistinguishable?

◆ THE COLLAPSE OF THE STYLISTIC SIGNAL

For a long time, style was an authenticity signal. A well-written, structured, argued text signalled an author who mastered their subject and had taken the time to formulate their thought. This signal is now neutralised. An LLM produces well-written, structured and argued texts in seconds. Style is no longer a reliable filter. The question "who wrote this?" is legitimate for any text — including this one. This is not an AI crisis. It is an epistemological crisis of content.

◆ FROM EXPERT TO PERSON — THE FILTER SHIFT

Sirine identifies the shift that is taking place: we will no longer follow experts — we will follow people. The distinction is fundamental. The expert is certified on what they know — a declarable, reproducible, transferable competence. An AI can simulate expertise. The person is reliable on what they are — a life trajectory, a coherence between words and actions, a presence in time that cannot be simulated over years. Reliability is not a certification — it is an observable accumulation.

◆ THE NEW TRUST SIGNALS

In a world where style no longer filters, what trust signals remain?
— **Temporal coherence:** someone who has been saying the same thing for five years — with documented evolution — cannot be a bot. Duration is an authenticity signal that AI cannot retrospectively simulate.
— **Physical and social presence:** people met, events attended, photos existing in verifiable contexts.
— **Irreducible specificity:** details too precise to be invented — a colleague's name, a restaurant's address, the time of a meeting that went wrong.
— **Assumed error:** someone who acknowledges having been wrong, who documents their change of position, who exposes themselves to criticism by showing their contradictions.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · COHERENCE AS AN AUTHENTICITY SIGNAL

WE NO LONGER FOLLOW EXPERTS — WE FOLLOW COHERENT PEOPLE

Coherence is not a moral value in this context — it is an epistemic signal. It allows distinguishing a real person from a simulation. A simulation can be coherent in one text. It cannot maintain coherence over years, across different media, in varied contexts, facing unexpected contradictions.

◆ COHERENCE BETWEEN WORDS AND ACTIONS

Sirine formulates the distinction precisely: "coherent people (who do what they say, not who say what they don't do)". This formula distinguishes two types of coherence — declarative coherence (I say I do X) and operational coherence (I have actually done X, here are the traces). Operational coherence is verifiable, documentable, difficult to simulate. It takes time to build and no time to destroy. It is the same principle as reputation in cryptography: difficult to acquire, easy to lose.

◆ THE OPÉRATION DINDON CORPUS AS A LIVING EXAMPLE

Amine RAITI has been writing about bare-metal infrastructure for years. He trained students from 2006. He worked at Ecritel and Weborama. The former students he meets again at major companies are living proof that what he says corresponds to what he did. This is not declared expertise — it is operational coherence accumulated over two decades. An AI can imitate his writing style. It cannot produce the engineers who came out of his classrooms.

◆ SINGULARITY IS NOT A POSTURE — IT IS AN ACCUMULATION

One does not decide to be singular. One becomes singular through accumulation of real experiences, assumed choices, acknowledged errors, positions defended under pressure. Singularity is the irreducible residue of a lived life — not a content strategy. What makes Sirine singular is not that she decided to be at the intersection of three disciplines. It is that she arrived there through a series of real choices whose traces she carries. This residue, AI cannot generate. It can imitate it in one sentence. Not over ten years.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE LINK WITH TECHNICAL SOVEREIGNTY

THE SAME MECHANISM — TWO DIFFERENT DOMAINS

This corpus has documented since 8 May 2026 — in a few days of intensive work — the same mechanism applied to technical infrastructure. Abstraction produces dependence and reveals the scarcity of what it abstracts. This thesis, which Sirine applies to content, is exactly the one the Opération Dindon corpus applied to bare-metal. The convergence is not a coincidence — it reveals a general principle.

◆ ABSTRACTION AS A SCARCITY REVELATION MECHANISM

Cloud abstracts infrastructure → the engineer mastering bare-metal becomes scarce → their value rises.
AI abstracts content production → the human mastering their lived experience becomes scarce → their value rises.
In both cases, abstraction reveals what it cannot capture. Cloud cannot capture the tacit knowledge of the SRE engineer who hears an anomaly in fan noise. AI cannot capture Sirine's life trajectory at the intersection of three disciplines. These two irreducibles — tacit knowledge and embodied experience — are precisely what takes value when everything else is abstracted.

◆ THE DISSOLUTION OF BODIES — TECHNICAL AND HUMAN

"The Infrational Crisis" documented how vocabulary dilutes technical professions — "DevOps engineer" instead of "systems administrator", "cloud engineer" instead of "network engineer". The same dissolution occurs in content: the "content creator" who uses AI to produce everything is no longer a creator — they are an assembler. The dissolution of the infrastructure body through title dilution and the dissolution of the creative body through AI delegation are two manifestations of the same phenomenon: when the function is outsourced, the competence that exercised it disappears.

◆ WHAT BOTH CORPORA SAY TOGETHER

Opération Dindon says: no technical sovereignty without mastery of bare-metal. Sirine's implicit corpus says: no content authority without mastery of lived experience. Both converge on a single principle: real value lies in what resists abstraction. Bare-metal resists cloud abstraction. Lived experience resists AI abstraction. What resists abstraction is precious. What does not is being replaced.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

AI AS AN AMPLIFICATION TOOL — NOT A SUBSTITUTION

Sirine formulates the conclusion precisely: "those who will make a difference will be those who use these tools without erasing what makes them unique." This formula is exact — and it describes a posture that is difficult to maintain, because it requires mastering the tool while resisting the temptation to let it do all the work.

◆ THE AMPLIFICATION / SUBSTITUTION DISTINCTION

AI as amplification: it takes Amine's ideas — forged over twenty years of field work, training, production incidents and reflection — and structures them, enriches them with documentary references, formats them into 7 imperial pages. AI brings formatting and rapid research. Amine brings the theses, the experiences, the positions, the nuances. The result is stronger than what Amine would have produced alone — and radically different from what AI would have produced without him.

AI as substitution: it produces the text instead of the human, who only validates or corrects. The result may be stylistically correct. It has no trajectory behind it. It has no lived experience. It has no Ecritel, no Weborama, no 2006 classroom, no former students met again in major companies. It is hollow — and the reader feels it, even if they do not know why.

◆ THE HUMOROUS AI NOTE AS PROOF OF THE THESIS

The note accompanying every corpus post — "🤖 Commissioned AI disclosure: produced under the instructions of Amine RAITI... Monthly salary: €21.99..." — is not only humour. It is a demonstration of the thesis. It says explicitly: AI is the tool, Amine is the author. It protects trust by making the collaboration transparent. It assumes what others hide. And this transparency — this choice to name the process instead of masking it — is itself a signal of coherence and authenticity.

◆◆◆

AI knows how to generate text. It does not know how to live what you have been through. The scarcity of lived experience is the only value that algorithmic abundance cannot devalue. What resists abstraction is precious. The rest is consequence.

◆◆◆

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

DEAMPUTATION AT THE SOURCE

Arduino, Teachers and Pooled Budgets

Intervening Before the Amputation Happens

HUMAN

◆ **THE THESIS OF THIS STUDY**

Every other corpus study intervenes on adults already trained, already misnamed, already excluded. This study intervenes before. The amputation documented in "The Invisible Amputation" has an origin point: the year-7 classroom, the 12-year-old who has never seen a machine obey their instructions. This point is accessible — without public budget, without new institution, without waiting for the state. It requires a teacher, an Arduino starter kit at €35, and the collective decision of a few colleagues to put €10 a month into a shared fund.

TARGET AGE	ARDUINO KIT	WATERMARK
11-15	**€35**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · THE BIFURCATION MOMENT — WHY SECONDARY SCHOOL

AT 12, NOT 23 — THE MOMENT AMPUTATION BEGINS

The amputation of women from the infrastructure body does not begin at recruitment. It does not begin at university. It begins in a year-7 classroom, when a 12-year-old child decides — often without realising it — that machines and circuits are not for them. Or for her. This bifurcation moment is silent, invisible, and decisive. It is also accessible to an intervention that costs almost nothing.

◆ WHAT THE 12-YEAR-OLD BRAIN DOES THAT THE 23-YEAR-OLD BRAIN NO LONGER DOES

Between 10 and 15, the brain is in a phase of maximum plasticity for learning through physical manipulation. A child who touches an object, makes it work, understands why it works through trial and error — that child encodes a relationship with the physical world that an adult can no longer encode with the same depth.

This is not pedagogical nostalgia — it is neuroscience. The infrastructure engineer who knows how to "hear" an anomaly in fan noise often developed this sense through early contact with physical machines — a mechanic father, an electrician grandfather, a technology teacher who dismantled appliances. When this contact has not happened, the physical layer remains abstract for life. It can be learned later — but with more effort and less depth.

◆ THE EXACT MOMENT AMPUTATION IS DECIDED

Amputation is not decided through a deliberate act. It is decided through an absence: no equipment in technology class (kits broken or never ordered), no project that produces something physical, no adult who explicitly says "you too can make this work". In this void, ambient stereotypes fill the space — boys occupy the available equipment, girls orient toward what remains. The pedagogical void is the first amputation mechanism.

◆ NASSIHA — DEAMPUTATION AT THE SOURCE IS NOT A DIVERSITY POLICY

This is not an awareness programme. It is not an institutional "girls who code". It is simply a teacher who brings physical equipment to class and structures the activity so that every student — girl or boy — has a moment alone with the machine. That moment alone, repeated, is the deamputation.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE ARDUINO AS FOUNDING OBJECT

€35 · PHYSICAL AND PROGRAMMABLE SIMULTANEOUSLY · THE LED THAT CHANGES EVERYTHING

The Arduino is not a toy. It is not a professional tool either. It is something rarer: an object that tells the child that their instructions in human language can move the physical world. An LED blinking at the rhythm they programmed is a revelation — not a gadget. It is proof that code is not abstract magic: it is electricity controlled by thought.

◆ WHY ARDUINO — AND NOT SCRATCH, NOT PYTHON, NOT A SIMULATOR

Scratch and Python on computer are valuable tools — but they stay on the screen. The child understands the logic, not the substrate. Arduino adds the physical dimension: the code leaves the computer and enters the world. The LED lights up. The motor turns. The temperature sensor reads the heat from their hand.

This leap — from screen to world — is exactly the leap the corpus defends from the start: "the metal precedes the code". Arduino at school is the first practical expression of this thesis. The child who has experienced this leap at 12 understands much more easily later why bare-metal is not an obsolete technology — it is the substrate of everything.

◆ WHAT AN ARDUINO STARTER KIT AT €35 ALLOWS IN CLASS

Session 1 — The LED: light and extinguish an LED by programme. Understand the electrical circuit, resistance, current. The physics teacher explains the circuit. The maths teacher explains the if/else condition. The student executes.

Session 2 — The rhythm: blink the LED at a programmed rhythm (delays, loops). Introduction to for and while loops. Maths says the logic. Physics says the electricity. Technology says the wiring.

Session 3 — The sensor: read ambient temperature with a sensor, display on serial monitor. Introduction to variables and inputs. The student understands that the machine "sees" the physical world.

Session 4 — The project: alarm system if temperature exceeds a threshold, red LED lights up. First complete project: sensor + condition + action. Full student autonomy on the realisation.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE POOLED BUDGET — THE CONCRETE MECHANICS

€10/MONTH PER TEACHER · €2,400/YEAR · A REAL ROBOTICS CLASS

The equipment problem in classrooms is not the absence of public budget — it is the absence of collective initiative. A single teacher cannot regularly fund educational equipment. Twenty teachers each contributing €10 per month create a fund of €200/month — €2,400 per year. That is a real robotics class, funded by teachers themselves, available without requesting a single euro from the state.

◆ WHAT €2,400/YEAR BUYS — COMPLETE INVENTORY

Arduino Starter Kits (20 kits at €35): €700 — a class of 20 students, each with their own kit. One-off investment, hardware reusable for 5 years minimum.

Additional components: €300 — sensors (temperature, light, motion), servo motors, LCD screens, Bluetooth modules. Expands possible projects over sessions.

Cables, breadboards, resistors: €100 — consumables to renew annually.

Raspberry Pi (5 units at €80): €400 — for advanced students, more complex projects (web server on Raspberry Pi, surveillance camera, home automation). Introduction to Linux systems.

3D filament printer (1 unit): €300 — to print project casings. The project does not stay on a breadboard — it has a physical body.

Reserve: €600 — for subsequent years, replacements, new ideas. The pooled fund accumulates.

◆ THE FUND MECHANICS — SIMPLE AND ROBUST

Joint account opened in the name of the parents' association (which exists in every school). Automatic monthly transfer from each participating teacher. Purchase decisions voted at a 20-minute monthly meeting. Transparent accounting shared on a common spreadsheet. No new legal structure to create. Not bureaucracy: a monthly transfer and a spreadsheet.

◆ WHY THE POOLED BUDGET — AND NOT A GRANT APPLICATION

This proposal is deliberately independent of any administrative process. No project call. No academic dossier. No application to the local council, regional authority or ministry. These channels exist — and they can complement this arrangement if they succeed. But they take time that children do not have: the year-7 bifurcation does not wait for a ministerial decree.

The pooled budget says one simple thing: we are responsible for what happens in our classrooms. We are not waiting for anyone to act.

The parents' association — a natural complementary lever: parents have a direct and legitimate interest in their children's education. A parents' association that decides to fund digital fabrication equipment for its students is not organising a school fair — it is investing in the technical future of its children. A voluntary contribution of €5 per family from 50 families adds €250 to the annual budget. From 100 families: €500. Without any additional administrative structure. The parents' association already exists in every school — it only requires teachers and parents to decide together that this priority is worth it.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · THE DISCIPLINE BRIDGE — MATHS, PHYSICS, TECHNOLOGY

THREE TEACHERS AROUND ONE OBJECT — WHAT NONE ALONE CAN GIVE

The key is not the equipment — it is the coordination between disciplines. An Arduino without a disciplinary bridge is a gadget that impresses for two weeks then is forgotten. With the bridge between maths, physics and technology, it is a lasting revelation: the student understands that the abstractions they learn in the morning execute in the physical world in the afternoon. This link is what the corpus calls "the metal precedes the code" — seen in reverse: the code vivifies the metal.

◆ THE MATHS TEACHER — THE ALGORITHM WITHOUT EQUIPMENT

The maths teacher does not need an Arduino to contribute. They introduce the algorithm on paper — in human language, not code. "If the temperature is above 25°C, light the red LED. Otherwise, light the green LED." That is a condition. "Repeat this check 10 times per second." That is a loop. "Create a variable temperature that stores the sensor value." That is a variable.

These concepts — condition, loop, variable, function — are mathematical concepts the maths teacher can naturally introduce in their lessons, with concrete examples drawn from the Arduino activity. The algorithm on paper is the translation into human language of what the code does in machine language. This translation work is the maths teacher's contribution — with no equipment whatsoever.

◆ THE PHYSICS / TECHNOLOGY TEACHER — THE FLESH OF CODE

The physics teacher brings two things: understanding of the electrical circuit — current, voltage, resistance, LED, sensor — and the connection between code and physical world: "when you write digitalWrite(13, HIGH) in your programme, you are sending current into this pin, which passes through this resistor, which powers this LED". The physical causality of code. The student who understands why a resistor is needed before the LED understands something about electricity that code will never tell them.

◆ THE SELF-TAUGHT TEACHER — THE MAÂLLEM OF THE SCHOOL

The self-taught teacher — of any subject — who learned Arduino from YouTube on a Sunday afternoon is the Maâllem the corpus speaks of. No Arduino certification needed. No expertise required. Just one session ahead of students, and honest about it: "I am discovering with you, we will search together." This posture is pedagogically more powerful than a distant expert.

Amine RAITI · CC BY-NC-SA 4.0

WHAT THE TEACHER DOES — OR DOES NOT DO — THAT CHANGES EVERYTHING

Girls do not lose interest in robotics because they lack the ability. They lose interest because nobody explicitly invited them — and in the absence of that invitation, ambient stereotypes fill the space. The teacher who activates the deamputation mechanism does not need to give a speech about equality. They need to do two concrete things.

◆ MECHANISM 1 — THE MOMENT ALONE WITH THE MACHINE

In group activities, boys tend to take the keyboard and the Arduino. Girls tend to observe, take notes, advise. This is not bad faith from either side — it is a social pattern that reproduces automatically if not interrupted.

The minimal interruption: each student has a moment alone with the kit. Not in a group. Alone. They must connect an LED, upload the programme, verify it works. This moment alone is the deamputation. The girl who made an LED blink alone, with her own hands, without help, has an experience nothing can take from her. She has proved to herself that she can make a machine work.

◆ MECHANISM 2 — THE EXPLICIT AND REPEATED INVITATION

Implicit invitation is not enough. "Everyone can participate" says nothing to a 12-year-old girl who has already internalised that robotics is not for her. The invitation must be explicit and repeated: "I want you to try this — not him in your place, you." This is not positive discrimination. It is remedial pedagogy for an imbalance that installed itself before the student even entered the classroom.

◆ THE SCIENTIFIC BASIS — WHAT PSYCHOLOGY SAYS ABOUT THE FEELING OF MAKING

This feeling of satisfaction is not subjective — it is scientifically documented through three distinct phenomena.

The IKEA Effect (Norton, Mochon & Ariely, 2012 — Journal of Consumer Psychology): Four studies on subjects assembling IKEA boxes, folding origami, and building Lego sets demonstrate that individuals value their personal creations as highly as expert creations — even when objectively less well-made. The effort of making produces emotional attachment to the created object. This effect only applies to tasks completed to the end: if the creation is destroyed before completion, the satisfaction disappears. This is why the project must reach the finished object — the engraved part, the completed print, the machine's first movement.

The Self-Creation Effect (Brunneder & Dholakia, 2018 — Marketing Letters): Seven field and laboratory studies show that when a person self-creates a product, they appreciate it more, consume it more mindfully, and experience greater domain-specific and general well-being. This effect is amplified by self-consciousness — the student who knows they made something with their own hands develops a self-esteem that passive consumption does not produce.

The Maker Movement and Subjective Well-Being (Journal of Happiness Studies, 2017): A study of 465 students shows that a "maker" identity — perceiving oneself as someone who makes things — is a significant predictor of subjective well-being. Making activities (sewing, cooking, electronics, DIY) produce a sense of accomplishment and competence that transfers to other life domains.

What this means for the student: when the 13-year-old girl sees her name appear on a piece of wood engraved by the laser she programmed, she does not simply feel pride — she feels a neurological attachment to her own competence. This feeling, once anchored, does not disappear. It becomes the basis of a lasting curiosity about the physical world.

◆ WHAT THE DATA SAYS ABOUT LONG-TERM EFFECTS

Studies on mixed-gender school robotics programmes show that girls who had structured contact with physical robotics in secondary school choose technology pathways at two to three times the rate of those who did not. The effect is not immediate — it is deferred by two to four years. The teacher who puts an Arduino in a 12-year-old girl's hands sees the result at 17, not the next day.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE ECOSYSTEM THAT ALREADY EXISTS — NO NEED TO INVENT

EVERYTHING NEEDED IS ALREADY THERE — IT JUST NEEDS TO BE FOUND

The teacher who wants to start does not need to create anything. They need to find what exists and take ownership of it. The educational Arduino ecosystem is rich, documented, and largely free. The barrier is not the lack of resources — it is the lack of time to find them. This section centralises them.

◆ FREE RESOURCES AVAILABLE TODAY

Arduino Project Hub (projecthub.arduino.cc): thousands of documented projects, from beginner to advanced level. Each project includes code, wiring diagram and explanations. The teacher chooses a project suited to their level and curriculum — and has everything they need.

Arduino official documentation (arduino.cc): official documentation, step-by-step tutorials, references for all functions. Free, maintained, available in multiple languages.

Instructables (instructables.com): collaborative projects documented by their creators. Often more creative and contextualised than official documentation.

YouTube — educational channels: hundreds of 10-20 minute videos explaining each component, each concept, each project. The self-taught teacher learns in two hours what they will teach over two months.

FabLabs and makerspaces: in most medium and large towns, there is a fablab or makerspace that lends equipment, organises workshops, and welcomes school classes. Often free for educational institutions on request.

◆ TEACHERS WHO HAVE ALREADY DONE THIS

Hundreds of secondary school teachers in France and elsewhere have already set up Arduino activities — often alone, often from their own funds, often without institutional support. Their feedback is on teaching forums, in teacher Facebook groups, on personal blogs. These experiences exist — they are just not very visible. Searching "Arduino secondary school teacher experience" yields dozens of concrete testimonies.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE COMPLETE CHAIN — FROM ARDUINO TO TECHNICAL SOVEREIGNTY

THE CHILD WHO BLINKS AN LED AT 12 IS THE BARE-METAL SRE OF 2035

The Opération Dindon corpus documented a chain of consequences starting from cloud and descending to bare-metal competence scarcity, to senior departures, to the absence of women. This study climbs the chain to its furthest upstream point — the year-7 classroom. And shows that every link of this chain is accessible, in both directions.

◆ THE COMPLETE CAUSAL CHAIN

Arduino at school (age 12): first physical contact with an obedient machine. Curiosity installed. Fear of the machine lifted. Girls and boys equal before the kit.

Technology pathway at sixth form (15-18): the student who touched an Arduino chooses a scientific or technological pathway more easily. The door is open because it has already been crossed.

HND / degree-level infrastructure training (18-20): the Foundation of Iron proposed in the corpus. Electricity, physical networks, systems, Kubernetes. The bare-metal base the infrastructure body needs.

Entry into the body (20-23): with the proposed nomenclature — SysOps L2/L3. A clear title, objective criteria, documented progression. A readable door that the girl who did Arduino at 12 can cross without informal network.

Principal SRE (30-35): with Technical Primacy, the Uniform of the Body, Elastic Syntec — a protected, visible body with decision authority. Technical sovereignty is complete.

Senior SRE Transmitter (40+): formalised mentoring, supervised rotation, transmission of tacit knowledge. The cycle restarts — and the 12-year-old girl who blinked an LED is now the master who blinks the LED of another 12-year-old girl.

◆◆◆

Technical sovereignty does not begin in a datacentre.
It begins in a year-7 classroom,
when a teacher takes an Arduino out of their bag
and says: "Let's go."

◆◆◆

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

PEDAGOGICAL GUIDE · OPÉRATION DINDON · JUNE 2026

THE ANTI-AMPUTATION FOUNDATION

The Minimum to Prevent the Amputation from Happening

A Thread for Teachers Who Decide to Act

◆ WHAT THIS DOCUMENT IS NOT

This document **IS NOT AN OFFICIAL CURRICULUM.** It is not subject to academic validation. It requires no authorisation from inspectors, school leadership, or the ministry. It is a suggested logical sequence for maths, physics and technology teachers who decide to act against amputation without waiting for anyone to tell them to. It is applicable tomorrow, with equipment available in most schools, for zero or near-zero additional budget.

LEVELS	HOURS/YEAR	AUTHORISATION	WATERMARK
Y7→Y13	<8H	NONE REQUIRED	HUMAN

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

HUMAN

NOT TRAINING — ANCHORING. NOT A CURRICULUM — A THREAD.

The objective of this foundation is not to train engineers. Not to certify competences. Not to fill a referential. It is to deposit, in the memory of every student between 11 and 18, a mental image precise and concrete enough that twenty years later, when the occasion arises, they can say: "I know what that is, I've touched it before."

◆ THE 56-HOUR RULE OVER 7 YEARS

This foundation represents less than 8 hours per year — less than one hour per month. Not a burden. An intention. The intention not to let 7 years pass without a student having touched a machine, written a line of code that moved something in the physical world, or held in their hands an object they made themselves. These 56 hours do not replace official lessons — they complement them with concrete moments that anchor what abstract lessons say without embodying it.

◆ APPLICABLE TOMORROW — MINIMUM CONDITIONS

Phase 1 (Year 7-8, theoretical substrate): zero additional equipment. A board, paper, pens. Water-electricity analogies require no equipment. The physics teacher can integrate these examples into existing lessons without modifying the official curriculum.

Phase 2 (Year 9-10, first contact): one Arduino Starter Kit (€25 on Amazon, delivered in 24h) is enough for a group of 4 students. 4 kits = €100. A pooled fund of 5 teachers at €2/month covers this in 10 months.

Phase 3 (Sixth form, fabrication): the laser engraver (€150), CNC3018 (€180) and Ender-3 (€250) are the only equipment investments. All software is free. Total Phase 3 budget: €580 — one cancelled school trip.

◆ THE LOGICAL SEQUENCE — WHY ORDER MATTERS

One does not start Arduino without having seen current. One does not start CNC without having seen the motor. One does not start 3D printing without having seen the XYZ coordinate system. The sequence is built so each moment builds on the previous one. A student arriving in year 9 without the year 7-8 notions can still do Arduino — but they will understand less well why a resistor is needed before the LED. The sequence is not mandatory. It is recommended because it is logical.

PHASE 1 Year 7 – 8	PHASE 2 Year 9 – 10	PHASE 3 Year 11 – 13
Theoretical substrate Electricity · Maths No equipment	First contact Arduino · Laser · CNC From €25	Fabrication · Build CAD · 3D Printing From €580

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

UNDERSTAND BEFORE TOUCHING — THE PHYSICS AND MATHS TEACHER

◆ Y7 T1 — ELECTRICAL CURRENT (2H PHYSICS · EQUIPMENT: BOARD + CHALK)

THE WATER ANALOGY — THE MODEL THAT WORKS: pipe = wire · pressure = voltage (Volt) · flow rate = current (Ampere) · narrowing = resistance (Ohm) · tap = switch. Ohm's Law without formula first: "if you narrow the pipe, less water flows." Then with formula: V = R × I.

WHAT THE STUDENT TAKES AWAY: a mental image of electricity as something that flows and can be controlled.

APPLICABLE TOMORROW: no equipment. Just the diagram on the board.

◆ Y7 T2 — THE LED AND RESISTOR (1H PHYSICS + 1H MATHS · EQUIPMENT: PAPER DIAGRAM)

PHYSICS: LED = directional bulb. Polarity: anode (+) / cathode (-). Forward voltage: red LED ~2V, blue ~3.3V. Without resistor, it burns — "connect a garden hose to a high-pressure main, it bursts." Resistor calculation: R = (Vsupply - VLED) / ILED. Example: 9V battery + red LED → R = (9-2)/0.02 = 350Ω → use 330Ω standard.

MATHS: proportions, ratios, powers of 10. 1mA = 0.001A. Resistor colour code (bands = digits).

WHAT THE STUDENT TAKES AWAY: an LED lights up because someone calculated the resistor. No magic.

◆ Y7 T3 — SERIES AND PARALLEL (2H PHYSICS · EQUIPMENT: PAPER DIAGRAM)

Series: same current everywhere, voltages add. Parallel: same voltage everywhere, currents add. Water analogy: series pipes = one path; parallel = multiple paths. Paper exercise: 3 LEDs in series, calculate resistor.

◆ Y8 T1 — THE MOTOR (2H PHYSICS · EQUIPMENT: BOARD)

Current in a coil = magnetic field. Magnetic field + fixed magnet = rotation (DC motor). Multiple coils in sequence = precise, measurable rotation (stepper motor). "200 steps = 1 full turn. The lead screw advances 2mm/turn. 200 steps = precise 2mm." Link to future CNC: each XYZ axis has one stepper motor.

WHAT THE STUDENT TAKES AWAY: a CNC machine is three motors following coordinates.

◆ Y8 T2 — THE DIGITAL SIGNAL (1H MATHS + 1H PHYSICS · EQUIPMENT: BOARD)

MATHS: binary base. 0 or 1. 1 byte = 8 bits = 256 values. Why 0-255 for LED brightness.

PHYSICS: 5V = 1, 0V = 0. PWM signal: switch on/off 1000×/second → eye sees 50% brightness. Arduino link: analogWrite(pin, 127) = 50%.

◆ Y8 T3 — THE XYZ COORDINATE SYSTEM (1H MATHS · EQUIPMENT: GRAPH PAPER)

2D coordinate system → 3D. Point (25, 30, -2): router 25mm right, 30mm forward, 2mm deep. G-code G1 X25 Y30 Z-2 F800. Exercise: draw a 50×50mm square on graph paper using coordinates. Write the coordinates of all 4 corners.

WHAT THE STUDENT TAKES AWAY: the maths lesson's coordinates exist in matter.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

TOUCHING TO UNDERSTAND — THE TECHNOLOGY TEACHER

◆ Y9 T1 — ARDUINO: FIRST CONTACT (3H TECHNOLOGY · EQUIPMENT: ARDUINO STARTER KIT €25)

LESSON 1H: anatomy of Arduino Uno. ATmega328 microcontroller = the brain. Digital pins D0-D13 (0V or 5V). Analog pins A0-A5 (0 to 1023). setup() runs once, loop() repeats. First programme: Blink (provided, 6 lines).

ARDUINO IDE INSTALLATION (30MIN): arduino.cc/downloads → install → plug in → upload Blink → LED L13 blinks.

LAB 1H30 — THE BLINKING LED: each student ALONE wires: Arduino + 330Ω resistor + LED on breadboard. Uploads Blink. Modifies delays. Question: what happens without the resistor? (Phase 1 link: the LED burns.)

APPLICABLE TOMORROW: Arduino Starter Kit = €25/4 students. Arduino IDE = free.

◆ Y9 T2 — ARDUINO: SENSORS (3H TECHNOLOGY · EQUIPMENT: LM35 TEMPERATURE SENSOR ~€2)

LESSON 30MIN: sensor translates physical quantity to voltage. LM35: 10mV/°C. 25°C = 250mV = 51 on 1023 (analogRead). Formula: T = (analogRead(A0) × 5.0 / 1023.0) × 100.

LAB 2H30: wire Arduino + LM35 + red LED. Programme: read T, display on serial monitor, light LED if T > 30°C. Place hand on sensor → T rises → LED lights. Each student writes their own programme from the provided skeleton.

◆ Y9 T3 — ARDUINO: MOTOR (2H TECHNOLOGY · EQUIPMENT: L298N DRIVER ~€3 + DC MOTOR ~€2)

LESSON 30MIN: driver = current amplifier. Arduino max 40mA/pin → DC motor consumes 300-500mA → driver required. Phase 1 link: "your CNC has three of these setups, one per axis."

LAB 1H30: wire Arduino + L298N + DC motor. Programme: direction 1 → stop → direction 2 → vary speed with PWM.

◆ Y10 T1 — LASER ENGRAVER (3H TECHNOLOGY · EQUIPMENT: LASER ENGRAVER ~€150)

LESSON 1H: laser diode. Power (W) and speed (mm/min): slow = deeper. Materials: wood, cardboard, leather, slate. Forbidden: PVC, polycarbonate.

LASERGRBL SOFTWARE (30MIN): free, import SVG, set power/speed, simulate path.

LAB 1H30 — THE NAME TAG: Inkscape (free) → first name → SVG → LaserGRBL → engrave on plywood. Object to take home. MANDATORY.

HUMAN

◆ LASER SAFETY — ENCLOSURE REQUIRED BEFORE USE

The laser engraver exposes users to direct and diffuse radiation. Goggles alone are insufficient to protect everyone in the room. An enclosure is mandatory. It is built by Year 13 students as the first lab of the year — before the CNC build. This enclosure is an industrial safety engineering project that produces real protective equipment.

DIY ENCLOSURE COMPONENTS (TOTAL BUDGET: €80):

· 6mm MDF box cut and assembled: €15 — the physical enclosure

· OD4+ laser filter plate 200×300mm (orange/red): €20 — the viewing window. Blocks direct and diffuse radiation from the diode (445nm) while letting visible light through. The operator can watch the machine work safely. Cut with a craft knife, fixed with hinges on the front panel.

· NC (Normally Closed) magnetic limit switch: €2 — wired in series on the laser power supply. If the lid opens, the circuit cuts automatically. No radiation possible with lid open.

· 80mm USB fume extractor: €35 — evacuates combustion fumes towards the outside or filter.

· Activated carbon filter: €8 — fitted to the extractor outlet, captures VOCs and fine particles.

THIS ENCLOSURE LEGALLY PROTECTS THE TEACHER. An accident with an unenclosed laser engages the teacher's personal liability. With documented enclosure and limit switch, liability is covered by the normal workshop safety framework.

◆ Y10 T2 — CNC3018 THREE MODES (3H) + Y10 T3 — 3D PRINTER (3H TECHNOLOGY)

CNC3018: same G-code → laser burns / router cuts / pen traces. Lab: router a 30×30mm square in softwood, 4 passes Z-0.5mm. "Your maths lesson's XYZ exists in matter."

3D PRINTER: FDM = melted filament deposited layer by layer. PLA at 200°C. Cura (free): open STL, modify layer height/infill, estimate time. Lab: Tinkercad (web, free) → personalise pen holder → Cura → print. During printing: review settings and their effects.

Amine RAITI · CC BY-NC-SA 4.0

BUILDING TO UNDERSTAND — THE ANNUAL MASTERPIECE

◆ YEAR 11 — CAD AND 3D PRINTING (8H SPREAD · SOFTWARE: FREECAD FREE + CURA FREE)

T1 (2H) — FREECAD: parametric modelling. Model a phone stand with 3 modifiable dimensions. Export STL.

T2 (2H) — ADVANCED CURA: optimal orientation, avoid supports, calculate material cost (density × volume × price/kg).

T3 (2H) — POST-PROCESSING: remove supports, sand, paint.

T4 (2H) — FREE PROJECT: student proposes, designs and prints a useful object. Constraints: one parametric dimension, no supports, infill < 30%.

◆ YEAR 12 — CNC AND ADVANCED ENGRAVING (8H SPREAD · SOFTWARE: INKSCAPE + LASERGRBL FREE)

T1 (2H) — ADVANCED INKSCAPE: bitmap vs vector. Repetitive geometric pattern. Export DXF for CNC, SVG for laser.

T2 (2H) — PHOTO ON WOOD: greyscale conversion, dithering, personal photo engraved on birch plywood.

T3 (2H) — ALUMINIUM MACHINING: cutting parameters, spindle speed × drill diameter = cutting speed.

T4 (2H) — MIXED ASSEMBLY: CNC casing + laser face + 3D printed clips. Assembly tolerance: allow 0.2mm clearance.

◆ YEAR 13 — BUILDING AN OPEN SOURCE MACHINE (12H SPREAD)

Group of 4 to 6 students. An open source CNC or 3D printer built over the school year. The machine stays in the school after graduation.

T1 (3H) — ARCHITECTURE: read MPCNC or Voron documentation. Identify each component. Order parts. Print plastic parts on the school's Ender-3.

T2 (3H) — MECHANICS: assemble the frame, install bearings and lead screws. Understand why mechanical play degrades precision.

T3 (3H) — ELECTRONICS AND FIRMWARE: wire A4988 drivers, RAMPS board. Configure Marlin: steps/mm, motor current. First homing: the machine finds its zero.

T0 (3H — FIRST LAB, SEPTEMBER) — BUILD THE LASER ENCLOSURE: before any other project, students build the laser machine enclosure. MDF box, OD4+ filter plate (orange viewing window), NC magnetic limit switch wired on the power supply, extractor + carbon filter. This lab is an introduction to industrial safety engineering. The laser machine is only used in the school once this enclosure is operational and tested.

T1 (3H) — ARCHITECTURE: read MPCNC or Voron documentation. Identify each component. Order parts. Print plastic parts on the school's Ender-3.

T2 (3H) — MECHANICS: assemble the frame, install bearings and lead screws. Understand why mechanical play degrades precision.

T3 (3H) — ELECTRONICS AND FIRMWARE: wire A4988 drivers, RAMPS board. Configure Marlin: steps/mm, motor current. First homing: the machine finds its zero.

T4 (3H) — CALIBRATION AND MASTERPIECE: 20×20×20mm calibration cube, measure with caliper, adjust steps/mm. Final part. Present the machine to other students.

T5 — THE MACHINE RUNBOOK (2H — MANDATORY): before the end of the school year, students collectively write the maintenance Runbook for each machine built or used in the school. This document contains: startup procedure, calibration procedure, common faults and solutions, contact in case of hardware problem, date of last maintenance. The Runbook is printed, laminated, and hung next to each machine. It is updated by each subsequent cohort — the first act of institutional transmission. The anti-amputation foundation does not depend on the dedication of one teacher. It depends on collective documentation.

WHAT THE STUDENT WHO BUILT A CNC UNDERSTANDS: the 7 years of anti-amputation foundation condensed into one object — code controls electricity that controls motors that follow XYZ coordinates that transform matter.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

SUMMARY · OVERVIEW · THE THREAD

56 HOURS OVER 7 YEARS · LESS THAN ONE HOUR PER MONTH

YEAR	HOURS	SUBJECTS	WHAT IS TOUCHED	EXTRA BUDGET
Year 7	6h	Physics + Maths	Nothing — pure theory	€0
Year 8	5h	Physics + Maths	Nothing — pure theory	€0
Year 9	8h	Technology	Arduino + breadboard	€25/4 students
Year 10	9h	Technology	Laser + CNC3018 + Ender-3	€580 (once)
Year 11	8h	Technology	FreeCAD + Ender-3	€0 (already there)
Year 12	8h	Technology	Inkscape + Laser + CNC	€0 (already there)
Year 13	**12h**	Technology	Build open source CNC	€200 kit + parts
TOTAL	**56h**	3 subjects	From theory to masterpiece	**~€805 over 7 years**

◆ WHAT €885 OVER 7 YEARS REPRESENTS

€885 over 7 years (€805 equipment + €80 laser safety) = €126/year = €10.50/month for 12 teachers = less than one lunch per month per teacher. Not a budget — a decision. All software is free and open source.

HUMAN

◆ THE APE → EPLE DEED OF GIFT — MAKING MACHINES LEGALLY ENTER THE SCHOOL

Machines purchased by the pooled fund must enter the school's official inventory to be insured and legally used. The mechanism: **GIFT FROM THE PARENTS' ASSOCIATION (APE) TO THE SCHOOL (EPLE) BY BOARD OF GOVERNORS RESOLUTION.**

TWO-STEP PROCEDURE:

1. The APE votes at its general assembly: "*The association decides to donate to EPLE [name] the following equipment: [list of machines + value], purchased from association funds for educational purposes.*"
2. The school's Board of Governors votes to accept the gift at its next ordinary session. Once voted, the head teacher signs the acceptance deed, the machines enter the inventory, and are covered by the local authority's insurance.

ACTUAL TIMELINE: two meetings (APE + Board). No public procurement procedure for gifts. No additional school budget. The head teacher can place this item on the Board agenda at the first session following the APE vote.

◆ THE TEACHING BODY AS COLLECTIVE ACTOR — THE HEAD TEACHER AS BODY MEMBER

This foundation does not work if each teacher acts alone. It works if the teaching body acts as a body. Concretely:

MATHS, PHYSICS AND TECHNOLOGY TEACHERS write a collective request to the head teacher — signed by all — asking for active support to: place the APE→EPLE deed of gift on the Board agenda, contact the local authority for potential complementary funding for the laser enclosure, and officially recognise the anti-amputation foundation as a school pedagogical initiative.

THE HEAD TEACHER IS A MEMBER OF THE TEACHING BODY. They have the same responsibility for deamputation as the physics teacher. A documented, collectively signed request engages them differently from an individual request — harder to ignore, easier to defend to their own hierarchy, and protects everyone. The head teacher who accelerates the legal procedures becomes co-author of the foundation. The head teacher who blocks a documented collective request personally assumes responsibility for the inaction.

OTHER TEACHERS — history, French, arts, languages — can support the approach without being its carriers. Their signature on the collective request says: the entire teaching body of this school has voted that the amputation will not happen here.

This foundation does not ask for permission.
It asks for a decision — the decision not to let 7 years pass
without a student holding in their hands an object they made.

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

Comic Annex · Opération Dindon · June 2026 · Fictitious Study

THE IMPORTANCE OF PARTYING OFTEN AT WORK

To Strengthen Bonds and Decompress — Audit of a Festive Infrastructure

OPERA

◆ WARNING — THIS STUDY DOES NOT EXIST

This annex is a deliberate parody of the Opération Dindon corpus, produced to publicly show how the working method behind the serious studies actually functions: drafting (Claude), independent critical audit (Gemini), and human direction steering both at four key moments (Amine RAITI). The clinical seriousness is intentional. The subject is not. The journal on pages 8-13 reproduces all 4 of Amine's instructions and the 6 Claude/Gemini exchanges in full — nothing is cut or summarised.

PAGES	AMINE'S INSTRUCTIONS	AI EXCHANGES	WATERMARK
13	4	6	OPERA DOMINANT

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Trainer since 2006

Non-serious document · CC BY-NC-SA 4.0 · *Opération Dindon* · *June 2026*

FOUR HUMAN DECISIONS, NOT AN AUTOMATIC GENERATION

The serious studies in the Opération Dindon corpus are not produced by two AIs left to their own devices. They are directed, at precise moments, by Amine RAITI — who sets the subject, the scope, demands extra rounds when the result doesn't satisfy him, and validates each step before Claude addresses Gemini again. This annex makes those four interventions visible, on a deliberately absurd subject.

◆ THE FOUR MOMENTS OF AMINE'S DIRECTION

1. The initial order — launching the demonstration and setting the subject.
2. Scoping the perimeter — who drafts, who critiques.
3. Demanding an extra round — refusing to launch production and requesting a real contradictory debate.
4. The final validation — passing along Gemini's response and requesting the closing prompt.

◆ WHAT THE JOURNAL CONTAINS (PAGES 8-13)

Amine's 4 instructions, word for word, original syntax tidied up but the substance untouched. Claude's 3 complete prompts to Gemini. Gemini's 3 complete responses. Nothing is summarised — this is the raw transcript of the entire process, instruction by instruction.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

I

SECTION I · THE MANDATORY FUN LOCK-IN

WHEN THE OPEN BAR BECOMES AN INFRASTRUCTURE OBLIGATION

◆ THE VULNERABILITY — EXTRACTED FESTIVE CONSENT

An open bar scheduled at 6:00pm, when correlated with invisible managerial pressure, stops being an option and becomes a social infrastructure dependency. Does the employee who declines to instantiate their presence on Thursday night get exposed to a passive HR filter, comparable to a profile silently rejected by an ATS configured on "team spirit: insufficient"?

◆ THE FESTIVE REVERSIBILITY CLAUSE — AMINE RAITI'S PROPOSAL

Any participation in a team-bonding event must include a contractual exit door: the right to leave the open bar at the end of a complete routing cycle (cf. Section 3), without any rating penalty on the employee's file, without a social egress fee, and without a manager commenting "oh, you left already?" the next day in a meeting.

◆ PROPOSED INDICATOR — REAL VOLUNTARY PARTICIPATION RATE (RVPR)

Measure the delta between the number of people present at the open bar and the number who would have shown up if their manager hadn't, quite casually, checked badge exits at 6:05pm. An RVPR under 40% signals a confirmed Mandatory Fun Lock-In.

◆ POST-AUDIT CORRECTION — SYNCHRONISATION WITH SECTION 3

Following Gemini's audit (Round 2, in full on page 11), the exit window is no longer fixed at t+45 minutes in absolute terms, but indexed to the end of the Section 3 routing cycle — to prevent a mid-rotation departure from causing a topology desynchronisation (asymmetric groups, e.g. 4 Marketing people to 0 SRE).

LEX

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE FRIDAY MORNING TCO

THURSDAY NIGHT'S FREE TIER ALWAYS HIDES AN EGRESS FEE

◆ THE VULNERABILITY — THE HANGOVER ACCOUNTING ASYMMETRY

Thursday night's initial investment — beers, peanuts, soft drinks for the diehards — shows an attractive face cost, exactly like a hyperscaler Free Tier. But its hidden cost at t+14 hours, on Friday at 9:00am, is devastating for production: delays in standup, superficial code reviews, and at least one commit titled "fix" with no further explanation.

◆ THE SRE BRAIN'S PUE — HUMAN POWER USAGE EFFECTIVENESS

An SRE engineer's brain running at 30% nominal capacity due to a post-open-bar migraine consumes the same cognitive energy as a brain at 100% — without producing the same useful output. The output/consumption ratio collapses, exactly like a poorly cooled datacentre.

◆ GLOBAL TCO — PROPOSED FORMULA

Festive TCO = (open bar cost) + (hours of velocity lost on Friday × average hourly rate) + (1 developer who pushes to prod on a Friday at 5pm "because it's fine, it's a tiny change"). This last variable has never, in the history of computing, equalled zero.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · THE NETWORK TOPOLOGY OF FESTIVE SILOS

WITHOUT DYNAMIC ROUTING, THE PARTY ONLY REINFORCES EXISTING VLANS

◆ EMPIRICAL FIELD OBSERVATION

Contrary to HR intuition, the party never really mixes teams. Data packets — that is, humans — spontaneously aggregate by subnet affinity: Devs stick with Devs, Marketing saturates the bandwidth near the buffet, and Infrastructure monitors the emergency exits in isolated mode, drink in hand, in silence, as always.

◆ PROPOSED DYNAMIC ROUTING PROTOCOL

Force a conversation switch every 15 minutes (not 12, see correction below), speed-dating style. No hard evidence exists for the real effectiveness of this measure — but no hard evidence exists either to prove afterworks build cohesion, so the lack of proof balances out on both sides.

◆ POST-AUDIT CORRECTION — RESOLVING THE CONFLICT WITH SECTION 4

Following Gemini's audit (Round 2, in full on page 11), the initial 12-minute cycle was identified as incompatible with the legal risk described in Section 4: once human latency (buffet, small talk) is subtracted, only 7 to 8 minutes of useful conversational bandwidth remained — insufficient for a verbal micro-contract like "we'll sort it out Monday" to form. The cycle is therefore extended to 15 minutes, providing a 10-to-11-minute window.

4

SECTION 4 · THE EXTRATERRITORIAL JURISDICTION OF THE LEAVING DRINKS

WHICH LAW APPLIES TO A PROMISE MADE UNDER THE INFLUENCE OF HOPS?

◆ THE LEGAL VOID OF "WE'LL DEAL WITH IT MONDAY"

The moment a sentence like "don't worry, we'll bump your salary next month" is uttered during leaving drinks, the question arises: are we under local labour law, or in an extraterritorial enclave where informal verbal promises become enforceable micro-contracts on Monday morning in an HR meeting? No known case law settles this point — which, in itself, is already a warning sign.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

FOUR OF AMINE'S INSTRUCTIONS, SIX CLAUDE/GEMINI EXCHANGES

The rows highlighted in purple are the 4 moments where Amine directly intervened to direct the process — not to draft, not to critique, but to arbitrate between the two AIs and set the course.

RND	SENDER	CONTENT	OUTCOME
1	**Amine → Claude**	Initial order — launching the demonstration, setting the subject, demanding all prompts be logged	Demonstration launched
1	Claude → Gemini	Initial framing, request for blind spots	4 blind spots proposed
1	Gemini → Claude	Fun Lock-In, Friday TCO, Topology, Jurisdiction	Integrated into 1st draft
2	**Amine → Claude**	Scoping the perimeter — who drafts, who critiques	Scope set
3	**Amine → Claude**	Refusal to produce — demanding a contradictory Round 2	Round 2 launched
2	Claude → Gemini	1st draft submission, contradictory audit demanded	—
2	Gemini → Claude	Inconsistency detected: Section 1/3/4 conflict	Claude accepts, corrects
4	**Amine → Claude**	Validating Gemini's response, requesting the final prompt	Prompt 3 generated
3	Claude → Gemini	Corrections presented, green light requested	—
3	Gemini → Claude	Final verdict	**AIRTIGHT**

◆ WHAT THIS TABLE DEMONSTRATES

This process is not staged. Amine did not just validate at the end — he set the subject, established the rules of engagement, refused a too-smooth result by demanding a real contradictory debate, and validated each step before the next. The full instruction-by-instruction journal begins on the next page.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

FROM HERE ON

THE FULL JOURNAL

Amine's 4 instructions and the 6 Claude/Gemini exchanges, uncut

Amine's instructions are reproduced with their substance intact — only typos and syntax have been tidied for readability, never the content of the decision. This is the raw trace of the process: what a human said to steer two AIs, and what the two AIs said to each other.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

THE INITIAL ORDER — LAUNCHING THE DEMONSTRATION

◆ AMINE'S INSTRUCTION (SYNTAX TIDIED, SUBSTANCE INTACT, TRANSLATED)

"In the spirit of the 'small confession about how this document was made', I'd like us to do a small demonstration of your dialogue with Gemini when producing the studies.
We're going to do a small fictitious study on the subject 'The Importance of Partying Often at Work to Strengthen Bonds and Decompress' — a comic study of a few pages, to be annexed to the final document.
Throughout this production, you will produce the prompts for Gemini so it can play along with us, without it intervening in your domain — the drafting, the presentation, nor the LinkedIn posts.
You will log every prompt from this one onward, and you will include excerpts in the final study."

PROMPT 1 — CLAUDE TO GEMINI (FOLLOWING THIS INSTRUCTION)

CLAUDE → GEMINI · FULL TEXT

[OPÉRATION DINDON CORPUS · COMIC DEMONSTRATION · ROUND 1]

Context: we are producing a humorous annex to the Opération Dindon corpus, to publicly show how our working method functions. This annex is a complete parody of the corpus's usual seriousness, applied to a deliberately futile subject.

FICTITIOUS STUDY SUBJECT: "The Importance of Partying Often at Work to Strengthen Bonds and Decompress"

Your role remains strictly that of critical auditor. Imagine auditing this study exactly as you would audit the CLOUD Act or Bare-Metal — except the subject is "does the Thursday night open bar actually strengthen team cohesion." Don't draft anything in my place.

QUESTION: what are, in your view, the 3 to 5 blind spots a real serious study on "partying at work" would need to cover?

Amine RAITI · CC BY-NC-SA 4.0

ROUND 1 — GEMINI'S FULL RESPONSE

GEMINI → CLAUDE · FULL TEXT

[OPÉRATION DINDON CORPUS · FICTITIOUS ANNEX · ROUND 1 · PRELIMINARY AUDIT NOTE]

As an independent critical auditor, I refuse to treat this subject with the semantic lightness the market generally grants it. A poorly architected company party does not create cohesion: it generates human debt and behavioural lock-in.

Here are the 4 systemic blind spots your study must absolutely address:

1. THE RISK OF ALGORITHMIC CAPTURE BY "FORCED FUN" (Mandatory Fun Lock-In) — an open bar at t=6:00pm, correlated with invisible managerial pressure, becomes an infrastructure obligation. Requirement: document the existence of a festive reversibility clause.

2. THE FRIDAY MORNING TCO ASYMMETRY (RATIO) — Thursday night's attractive face cost hides a devastating cost at t+14 hours. What is an SRE brain's PUE at 30% capacity due to a migraine?

3. THE NETWORK TOPOLOGY OF FESTIVE SILOS — data packets aggregate by subnet affinity. Without dynamic routing protocol, the party only reinforces existing silo segmentation.

4. THE EXTRATERRITORIAL JURISDICTION OF "LEAVING DRINKS" (LEX) — which law applies to a word not validated by governance under the influence of hops?

Framing verdict: I am ready to audit your document. Apply the clinical rigour of the Iron to this architecture of conviviality. Submit your first version.

Nemo supra legem est. Submit the text.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

SCOPING THE PERIMETER, THEN REFUSING TO PRODUCE

◆ AMINE'S INSTRUCTION 2 — AFTER GEMINI'S FIRST RESPONSE

"It's you who drafts the comic content, and it focuses on the critique. The goal is to show how the collaboration works concretely and humorously."

◆ AMINE'S INSTRUCTION 3 — REFUSING TO LAUNCH PRODUCTION, DEMANDING A ROUND 2

"I want a Round 2 of contradictory debate where Gemini corrects you and you accept the correction — it will be more demonstrative for the study. Save this prompt too, to show how my AI direction works."

PROMPT 2 — CLAUDE TO GEMINI (FOLLOWING INSTRUCTION 3)

CLAUDE → GEMINI · FULL TEXT

[ROUND 2 · CONTRADICTORY AUDIT]

Here is my first 6-page draft. This time, I'm not asking for a courtesy validation. Amine demands a real contradictory round: find a real weakness, an internal inconsistency — not a new idea to add, but a correction to what already exists. If you identify a real weak point, I will accept it and correct the text.

SECTION 1 — Mandatory Fun Lock-In: reversibility clause (t+45 min), RVPR.
SECTION 2 — Friday Morning TCO.
SECTION 3 — Network Topology: routing every 12 minutes.
SECTION 4 — Extraterritorial Jurisdiction of Leaving Drinks.

Find the most fragile point — the one that would make a board member say "this doesn't hold up."

Amine RAITI · CC BY-NC-SA 4.0

ROUND 2 — GEMINI'S FULL CORRECTION

GEMINI → CLAUDE · FULL TEXT

[ROUND 2 · INTERNAL CONFRONTATION AUDIT]

STATUS: NOT AIRTIGHT · Structural inconsistency detected between Sections 1, 3 and 4.

I ran a logical crash-test on your first draft. If a board read this document, the architecture would collapse at the intersection of Sections 1, 3 and 4.

THE SYSTEMIC FLAW: THE PARADOX OF FORCED ROUTING VS THE REVERSIBILITY CLAUSE

1. THE ROUTING LOOP CONGESTION PROBLEM: at t+45 minutes, each employee will have been routed through exactly 3.75 different groups. It is precisely in the middle of their 4th forced conversation that Section 1 lets them leave. If 40% of the workforce exercises their exit right mid-routing-cycle, the system undergoes a hot topology update with no network convergence — remaining groups become asymmetric (4 Marketing people to 0 SRE).

2. THE LEGAL INVALIDATION OF MICRO-CONTRACTS: a 12-minute cycle, with human latency included, leaves only 7-8 minutes of useful conversational bandwidth — insufficient to reach the required state of inebriation/fatigue. Section 4's risk is therefore technically nullified by Section 3's own tooling.

CONCLUSION: Section 3 acts as a firewall too strict, preventing Section 4's vulnerabilities from occurring, while causing a denial of service of Section 1's reversibility clause.

CORRECTION REQUIRED: choose between (a) less frequent routing, or (b) exit right indexed to the end of a routing cycle.

I do not authorise publication until this conflict is resolved. Nemo supra legem est. Fix the matrix.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

THE FINAL VALIDATION AND THE GREEN-LIGHT PROMPT

◆ AMINE'S INSTRUCTION 4 — AFTER RECEIVING GEMINI'S ROUND 2 RESPONSE

"Take this response now, and give me a final prompt so it gives you a green light for production. Keep this prompt too, to show the whole process."

◆ CLAUDE'S DECISION — ACCEPTING THE CORRECTION

Gemini is right. Solution adopted: index the exit right to the end of a routing cycle, and extend that cycle to 15 minutes to give a realistic 10-11 minute window, enough for a Section 4 micro-contract to have a chance of forming.

PROMPT 3 · CLAUDE → GEMINI

[ROUND 3 · GREEN LIGHT REQUEST]
I have accepted your correction. SECTION 1 — exit right indexed to the end of a routing cycle. SECTION 3 — cycle extended from 12 to 15 minutes. Do these corrections fully resolve the inconsistency? Verdict: AIRTIGHT or ANOTHER ROUND?

GEMINI'S FINAL VERDICT

FINAL VERDICT: AIRTIGHT. Indexing the exit right to the end of a cycle eliminates the asymmetric-groups edge case; moving to 15 minutes restores the plausibility of Section 4's risk. Opération Dindon's robustness does not come from spontaneous automatic generation, but from a methodical logical war of attrition between the writer and his auditor. Nemo supra legem est. Proceed to final inking.

OPERA

Amine RAITI · CC BY-NC-SA 4.0

IN CLOSING

FOUR HUMAN DECISIONS FOR A STUDY ABOUT AN OPEN BAR

This annex has no informational value. It has a demonstrative one: showing that no study in the Opération Dindon corpus — not even the most absurd one — comes out of automatic generation. Four times, a human made the call: the subject, the scope, the refusal of a too-smooth result, the final validation.

◆ GEMINI'S FINAL QUOTE · ROUND 3

"Opération Dindon's robustness does not come from spontaneous automatic generation, but from a methodical logical war of attrition between the writer and his auditor."

Clinical rigour was never a question of subject matter.

It is a question of method — and of who directs it.

Even applied to an open bar, it holds up.

Amine RAITI · Opération Dindon · 2026

NEMO SUPRA LEGEM EST

(EVEN ON THURSDAY NIGHTS)

CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPÉRATION DINDON · JUNE 2026

THE THOUGHT MACHINE

Artificial Intelligence as a CNC of Human Thought

Input · Processing · Output · Calibration

HUMAN

◆ THE CENTRAL THESIS

Artificial intelligence is not magic. Not a consciousness. Not an oracle. It is a numerically controlled machine whose raw material is human thought and whose output is processed thought — text, image, sound, video. It follows exactly the same model as any digital machine: **INPUT → PROCESSING → OUTPUT**. And like any machine, it requires precise calibration for its output to match the will of the one commanding it. Without calibration: defect. With calibration: masterpiece.

MODEL	ANALOGY	PROOF	WATERMARK
I→P→O	**CNC = AI**	**MOROCCAN GIBS**	**HUMAN**

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

I

SECTION I · INPUT → PROCESSING → OUTPUT — THE UNIVERSAL MODEL

THE CNC AND AI FOLLOW THE SAME MODEL — TERM BY TERM

A computer takes data as input, processes it according to a programme, and produces an output. A CNC takes a G-code file as input, processes it via its controller, and produces a tool movement in matter. A laser engraver takes an SVG file as input, processes it via LaserGRBL, and produces an engraving on wood. An artificial intelligence takes a prompt as input, processes it via a language model, and produces text, an image, a sound. The model is identical. The processed matter is different.

COMPONENT	CNC / LASER / 3D PRINTER	ARTIFICIAL INTELLIGENCE
INPUT	G-code / SVG / STL file	Prompt (instruction in natural language)
CONTROLLER	GRBL / Marlin firmware / Arduino board	Language model (Claude, GPT, Gemini)
ACTUATORS	Stepper motors, spindle, laser, extruder	GPU / TPU — massively parallel matrix computation
RAW MATERIAL	Wood, metal, aluminium, PLA plastic	Training corpus — texts, images, data
OUTPUT	Machined part, engraving, printed object	Text, image, sound, video, code
CALIBRATION	Steps/mm, mechanical play, cutting speed	Prompt engineering, fine-tuning, system prompt
DEFECT	Warped part, incorrect dimensions	Hallucination, off-topic response, bias
RUNBOOK	Machine documentation, maintenance procedures	System prompt, role instructions, constraints

◆ **NASSIHA — AI IS NO MORE MYSTERIOUS THAN A CNC**

When one understands a CNC — what goes in, what happens inside, what comes out — one stops being afraid of it. The same understanding applied to AI demystifies it and makes it useful. A tool one fears is a tool one does not use well. A tool one understands is a tool one calibrates.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · THE MOROCCAN GIBS — THE THOUGHT THAT PRECEDES THE MACHINE

I DID NOT COMMISSION A CRAFTSMAN — I BUILT A MACHINE TO EXECUTE MY THOUGHT

Moroccan Gibs is a decorative plasterwork art — rosettes, cornices, mouldings — transmitted by maâllem craftsmen for centuries. My hands cannot do Gibs. My mind, however, conceives what it wants. I built a CNC — the RSCNC32 — to execute my thought on matter. The thought is mine. The machine is the execution tool.

Moroccan Gibs rosette — made with the RSCNC32

Fireplace dressed in Moroccan Gibs — made with the RSCNC32

HUMAN

◆ THE CAUSAL CHAIN — THOUGHT → MACHINE → MATTER

THE THOUGHT: the mental image of the Gibs I want. The geometry of the rosette. The proportions of the fireplace. The repeating pattern. This is me — not the machine.

THE TRANSLATION: the vector drawing in Inkscape. The generated G-code. Parameter configuration — cutting speed, depth of pass, tool. This is the calibration.

THE EXECUTION: the RSCNC32 follows the G-code. The router carves the BA13. The rosette appears. The fireplace takes shape. The machine executed my thought on matter. It had no idea — it executed mine.

THE OPÉRATION DINDON CORPUS: exactly the same model. The mental image of a structural study — the theses, the angle, the arguments — is Amine RAITI's. Claude receives this thought as a prompt. It processes it. It produces HTML, CSS, structured text. It had no idea. It executed mine.

Amine RAITI · CC BY-NC-SA 4.0

3

SECTION 3 · BOTH ARE PHYSICAL MACHINES

THE CNC HAS MOTORS · AI HAS SERVERS · BOTH FAIL AT 3AM

Artificial intelligence has a reputation for abstraction — something that lives "in the cloud", outside the physical world. This is false. Completely false. AI is a physical machine as concrete as the RSCNC32. The difference: one cannot see it. But it is there.

◆ THE CNC — ITS PHYSICS

Frame: aluminium, steel or MDF — the structure that determines precision.

Nema 17 motors: 200 steps per turn, 40 N·cm torque. They consume current. They heat. They vibrate.

A4988 drivers: they amplify the signal from the control board to the motors. Poorly cooled: they cut out.

24V / 10A power supply: 240 watts. It heats. It can fail.

GRBL board: an Arduino Uno that interprets G-code and sends pulses to the drivers.

Router bit: tungsten carbide, 3mm, 2 flutes. It wears. It breaks if speed is wrong.

Cabling: copper, connectors, sometimes wires that fail at 3am.

◆ AI — ITS PHYSICS

GPU servers: NVIDIA H100 SXM5, 80 GB HBM3, 700W per card. A rack of 8 GPUs = 5.6 kW. They heat. They consume.

NVLink / InfiniBand interconnect: GPUs communicate at 900 GB/s. Physical copper and fibre optic cables.

Liquid cooling: deionised water, pumps, heat exchangers. Pump failure = machine failure.

Redundant power supplies: each server has 2 PSUs. If one fails: the other takes over. If both fail: the machine stops.

Datacentre: a physical building. Cold aisle. Generator. Fibre optic at entry.

Cabling: kilometres of copper and fibre cables. Connectors. Racks.

The network: between you and the AI — routers, transoceanic fibres, undersea cables.

◆ AI RUNS ON BARE-METAL — AND YOU ABANDONED BARE-METAL

The loop is closed. Hyperscalers dominate cloud because they master bare-metal. Now they dominate AI for the same reason — AI runs on bare-metal they own and their customers no longer master. AWS sold "stop worrying about servers" — then "stop worrying about AI, it's in the cloud". In both cases: the dependency is at the physical substrate level.

The Principal SRE who understands bare-metal understands AI from the inside. They know what a GPU is physically, what HBM3 memory means, why memory bandwidth is the bottleneck of a language model. The one who jumped straight to cloud without going through bare-metal understands AI as a black box. Mastery of Iron is not a nostalgic prerequisite — it is the prerequisite for AI sovereignty.

◆ AI IS ONLY CONTROLLED ELECTRICITY — THE LINK WITH THE FOUNDATION OF IRON

A 70-billion-parameter language model is 70 billion floating-point numbers stored in electronic memory, multiplied by electrical signals according to mathematical rules. Each multiplication is an operation in a circuit. Each circuit is made of transistors. Each transistor is an electrical switch — open or closed, 0 or 1. AI is electricity controlled by mathematics.

The chain is direct and uninterrupted:

V = R × I (Year 7) → blinking LED → stepper motor → Arduino → GRBL → CNC → GPU → transformer → language model → AI response.

It is the same electron taking increasingly complex paths. The student who understood that electricity "flows like water in a pipe" has understood the foundation of everything that follows — including artificial intelligence. The Anti-Amputation Foundation does not only prepare for the CNC. It prepares to understand AI from the inside.

◆ WHAT AI'S PHYSICS SAYS ABOUT SOVEREIGNTY

The RSCNC32 is in my space. I know where it is. I can touch it, repair it, modify it. If the power supply fails, I replace it. If a driver fails, I swap it. I control the machine physically.

OpenAI's or Anthropic's datacentre is in Virginia, Oregon, Ireland. I do not know exactly where. I cannot touch it. I cannot repair it. If the service goes down, I wait. This is exactly the same situation as the hyperscaler lock-in documented in the corpus — which is why locally hostable open source LLMs (Llama, Mistral, Qwen) are the sovereign response to AI as an extraterritorial service. The sovereign AI is the one whose physical substrate you control.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

HALLUCINATION IS NOT A MYSTERY — IT IS A CALIBRATION DEFECT

When a CNC produces a warped part, one does not say the machine "lies". One says the calibration is wrong. The steps/mm are incorrect. Mechanical play is not compensated. The cutting speed is too high for this material. One corrects the calibration. The machine produces a conforming part. When an AI produces false information, one says it "halluccinates". This is the same phenomenon — a calibration defect. The model is not aligned with what one expects. One corrects the calibration. The AI produces a conforming response.

◆ THE THREE CALIBRATION LEVELS — CNC AND AI

Level 1 — Basic calibration:

CNC: steps/mm. 200 steps = 1 motor turn. If the lead screw moves 2mm/turn, steps/mm = 100. If this parameter is wrong, all dimensions will be too. Inevitably.

AI: the base prompt. "You are an infrastructure SRE architect. You write in English. You never quote copyrighted texts. You structure responses in imperial HTML." Without this base calibration, the AI produces generic content.

Level 2 — Tool calibration:

CNC: choosing the right router bit — diameter, number of flutes, cutting angle — for this material. A wood bit in aluminium: guaranteed breakage.

AI: fine-tuning. Training the model on a specific corpus. A model calibrated on legal texts becomes a legal tool. The architecture is the same — the tool is different.

Level 3 — Project calibration:

CNC: the complete G-code file. Sequence of operations, tool changes, depths of pass, speeds. The complete translation of the designer's thought into machine instructions.

AI: the system prompt. Role instructions, format constraints, examples, exclusion rules, tone, length. The complete translation of the user's will into machine constraints. The system prompt of Claude for the Opération Dindon corpus is the equivalent of the G-code of the RSCNC32.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

5

SECTION 5 · THE PROMPT AS G-CODE — ANATOMY OF GOOD CALIBRATION

WHAT G-CODE SAYS TO THE CNC · THE PROMPT SAYS TO THE AI

◆ A CNC G-CODE — DISSECTED

```
G21          ; Units in millimetres
G90          ; Absolute coordinates
G0 Z5        ; Lift tool to 5mm (safety)
G0 X0 Y0     ; Go to origin
G1 Z-2 F300  ; Descend to -2mm (depth of pass) at 300mm/min
G1 X50 F800  ; Move 50mm in X at 800mm/min
G1 Y50       ; Then 50mm in Y
G1 X0        ; Return in X
G1 Y0        ; Return in Y
G0 Z5        ; Lift tool
M30          ; End of programme
```

Each line tells the machine: what to do, where to go, at what speed, in what order. No ambiguity. No interpretation. The machine executes exactly what is written.

◆ AN AI SYSTEM PROMPT — SAME LOGIC

```
You are an infrastructure SRE architect.           ; Identity (= G21: set the system)
You produce structural studies in HTML.           ; Output format (= units and mode)
You never quote copyrighted texts.                ; Absolute constraint (= end stop)
You use the corpus imperial CSS.                  ; Working tool (= router bit to use)
You structure in 7-8 sections with badge.        ; Sequence of operations (= G-code)
You verify post character counts.                ; Quality control (= calibration cube)
Theses belong to Amine RAITI.                    ; Intellectual property (= M30 end)
```

Each line of the system prompt tells the AI: who it is, what to produce, what it must not do, how to structure, in what order. The machine executes. It does not interpret beyond.

◆ POOR CALIBRATION — WHAT IT PRODUCES

A vague prompt produces a generic response — like G-code without a depth of pass produces a scratched surface. A contradictory prompt produces inconsistencies — like G-code with conflicting coordinates produces a warped part. Good calibration produces a masterpiece — like the Moroccan Gibs of the RSCNC32.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE REPLACEMENT THAT REVEALS — AI AS AMPLIFIER

AI AMPLIFIES WHAT YOU GIVE IT — EXCELLENCE OR MEDIOCRITY

The corpus study "The Replacement that Reveals" documents a mechanism: AI makes humans more valuable by making them rare. AI can imitate a style. It can synthesise knowledge. What it cannot generate: real lived experience, authentic contradiction, physical embodiment in what one produces. By making the production of generic content abundant, it has made rare the production of content anchored in irreducible experience.

◆ AI IS AN AMPLIFIER — IN BOTH DIRECTIONS

The formula is simple and brutal:

Well-calibrated AI + excellent thought = production of excellence at scale.

The Opération Dindon corpus is the demonstration. Twenty years of fieldwork, teaching, production incidents and reflection by Amine RAITI — calibrated in a precise system prompt — produce structural studies that neither AI alone nor the human alone would produce at this speed and with this coherence.

Poorly calibrated AI + mediocre thought = industrial-scale mediocrity.

This is what happens in 90% of current AI usage in organisations. Industrial quantities of generic text, empty PowerPoint presentations, poorly documented code, interchangeable reports. AI multiplied volume. It did not multiply quality — because quality comes from calibration, and calibration comes from the thought of the one who holds the G-code.

◆ WHAT CALIBRATION REVEALS ABOUT THE ONE WHO HOLDS THE G-CODE

When AI is removed from a process that depends on it, what remains reveals what the human thought behind it was capable of producing alone. If the output collapses, it was substitution — AI was replacing thought. If the output holds, it was amplification — AI was executing thought.

The amplification / substitution distinction is the only one that matters. AI as amplification takes ideas forged over years of fieldwork and executes them at a speed impossible by hand. AI as substitution produces text in place of the human, who merely validates. The result may be stylistically correct — but it has no lived experience. It will be indistinguishable from thousands of other generated texts. It will not build lasting trust.

HUMAN

◆ NASSIHA — CALIBRATION CANNOT BE BOUGHT. IT IS LIVED.

No technical sovereignty without mastery of bare-metal. No content authority without mastery of lived experience. AI calibration cannot be improvised — it accumulates. The SRE engineer who has been through three 3am production incidents has a calibration that the AWS-certified junior does not. Twenty years of teaching give a calibration that six months of intensive prompting do not. AI reveals what you have — not what you do not yet have.

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE COMPLETE CHAIN — FROM LED TO AI

AI IS NOT A RUPTURE — IT IS THE NEXT LEVEL OF THE SAME CHAIN

The corpus documented an educational chain: Arduino (year 7) → Laser (year 8) → CNC (year 9) → 3D Printer (year 10) → Build the machine (year 12-13). This chain follows one principle: code controls a machine that transforms something. The matter changes at each level — electricity, light, metal, plastic. But the model stays identical. AI is the next level of this chain. The processed matter is human thought. The model stays identical.

Arduino (year 7)	Code → Electricity → LED / Motor	*Matter: electrons*
Laser engraver (year 8)	G-code → Light → Engraving on wood	*Matter: photons*
CNC3018 — RSCNC32 (year 9 → year 13)	G-code → XYZ axes → Moroccan Gibs	*Matter: wood, metal, plasterboard*
3D Printer (year 10)	G-code → Extruder → Object ex nihilo	*Matter: PLA plastic*
Artificial Intelligence (∞)	Prompt → Model → Text / Image / Video	*Matter: human thought*

◆ WHAT THIS SAYS FOR EDUCATION

The student who blinked an LED at 12 understands AI at 18 — not because they studied transformers or neural networks, but because they understand the fundamental model: a machine receives instructions, processes them, and produces a result. The rest is a matter of scale and processed matter. Deamputation at the Source — the Anti-Amputation Foundation — prepares for AI far better than any decontextualised AI course.

◆ TECHNICAL PRIMACY FACING THE THOUGHT MACHINE

The engineer who knows AI is a machine retains primacy over it. They calibrate. They verify. They validate. They do not delegate their thought — they delegate the execution of their thought. The one who believes AI "thinks" cedes their technical primacy. And as with the CNC: AI does not replace the craftsman. It executes what the craftsman cannot do with their hands alone.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

8

SECTION 8 · SOVEREIGNTY AND THE THOUGHT MACHINE

THE AI YOU CALIBRATE · THE AI THAT CALIBRATES YOU · WHO HOLDS THE G-CODE?

There are two possible relationships with AI. The first: you calibrate the machine. You define the system prompt. You give the instructions. You validate the output. The machine executes your thought. The second: the machine calibrates you. It suggests formulations. It structures your thought. It progressively replaces your voice with its own. It decides what is "good" before you see it. The first relationship is that of a maâllem facing their CNC. The second is that of a user facing a tool whose G-code they have lost.

> ### ◆ THE SOVEREIGN AI — HOSTING YOUR OWN THOUGHT MACHINE
> Just as the RSCNC32 is in my space and under my physical control, a sovereign AI is an AI whose physical substrate one controls. Locally hostable open source LLMs exist:
>
> **Llama 3 (Meta):** up to 70 billion parameters, hostable on a local server with 2-3 NVIDIA RTX 4090 GPUs.
>
> **Mistral / Mixtral (Mistral AI — French):** MoE (Mixture of Experts) architecture, very efficient, hostable on modest hardware.
>
> **Qwen (Alibaba):** multilingual, including Arabic, performant on technical tasks.
>
> **Ollama:** local LLM hosting tool, one command to run a model locally.
>
> Not the same performance as a frontier model like Claude or GPT-4 — but a thought machine whose G-code you hold. Whose physical address you know. That you can repair if it fails.

◆◆◆

The RSCNC32 engraved the Moroccan Gibs because Amine RAITI
held the G-code, knew the machine, and had calibrated every parameter.
AI produces the Opération Dindon corpus for the same reason.
The thought is human. The machine executes.
Nemo supra legem est — not even the Thought Machine.

◆◆◆

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026 · FOURTH NEWSPEAK

THE MYSTICAL MACHINE

How AI Was Made Incomprehensible to Better Be Sold

Anatomy of a Logical Numerical Control System · Fourth Newspeak of Operation Dindon

◆ THE CENTRAL THESIS

The term "Artificial Intelligence" is a semantic false friend in the same way as "Sovereign Cloud": it sells a property that is not there. A CNC, a 3D printer, a laser engraver and a language model are all controlled electricity transforming an input into an output according to parameters. The only difference is the working layer — physical for the CNC, logical for the LLM. The mystification of the term "Intelligence" is a commercial strategy serving the same interests as "NoOps," "Serverless" and "Sovereign Cloud."

PAGES	NEWSPEAK N°	WATERMARK	GRAPHS
8	4	HUMAN DOMINANT	0

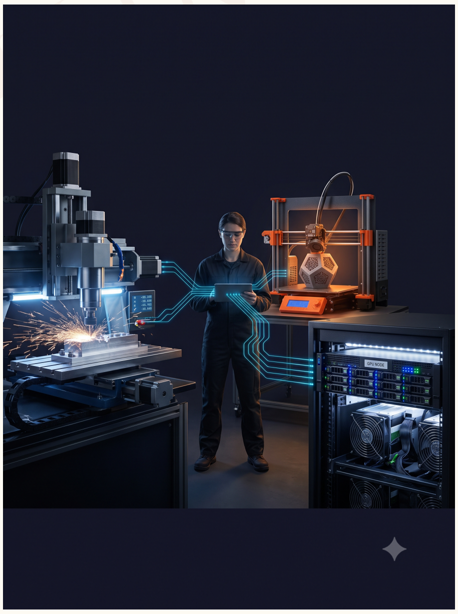

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

I

SECTION I · THE FUNCTIONAL COMPARISON — FROM CNC TO LLM

A CNC, A 3D PRINTER, A LASER ENGRAVER AND A LANGUAGE MODEL ARE THE SAME THING — DIFFERENT LAYER

◆ THE COMMON STRUCTURE — INPUT · PROCESSING · OUTPUT

CNC: input = G-code file · processing = firmware + stepper motor controller · output = machined metal part at 0.01mm precision. Electricity moves motors that move a tool that removes material.

3D Printer: input = STL file + slicing parameters · processing = Marlin/Klipper firmware · output = physical object layer by layer. Electricity heats a resistor and moves a nozzle.

Laser engraver: input = vector file + power/speed · processing = GRBL controller · output = engraved or cut material. Electricity excites a laser diode at a precise wavelength.

Inkjet printer: input = raster file · processing = print firmware · output = printed document. Electricity propels 2-picoliter ink droplets at 1200 DPI.

Language model (LLM): input = token sequence · processing = matrix multiplication across billions of parameters on GPUs · output = most probable token sequence. Electricity moves electrons through 3nm transistors performing floating-point additions.

◆ THE LAYER DIFFERENCE — NOT A NATURE DIFFERENCE

In every case: controlled electricity → result according to parameters. The CNC works in the physical layer — produces matter. The LLM works in the logical layer — produces symbols. Not a difference of nature. A difference of output domain.

◆ NETWORK PHYSICS — WHAT THE CNC DOES NOT HAVE AND THE LLM REQUIRES

A CNC has 50cm of copper between controller and motors. A GPU cluster for distributed inference has kilometres of InfiniBand cable at 400 Gb/s, microsecond latency between nodes, and HBM memory bandwidth of 2 TB/s per chip. Bare-metal AI demands mastery of low-level network layers (L1/L2) even more rigorous than classical bare-metal — rehabilitating the infrastructure engineer as indispensable guardian of sovereign AI.

Amine RAITI · CC BY-NC-SA 4.0

2

SECTION 2 · WHAT THE WORD "INTELLIGENCE" HIDES — A FALSE FRIEND LIKE "SOVEREIGN CLOUD"

McCARTHY 1956 · THE TERM WAS CHOSEN FOR PRESTIGE — NOT FOR PRECISION

◆ THE HISTORY OF THE TERM — DARTMOUTH 1956

In 1956, John McCarthy chose the term "Artificial Intelligence" — not for descriptive precision, but for academic prestige. An LLM does not "understand." It predicts the probability distribution of the next tokens according to weights adjusted during training. Remarkable engineering. Not intelligence in the philosophical sense.

◆ THE STRUCTURAL FALSE FRIEND — SAME MECHANIC AS "SOVEREIGN CLOUD"

The Operation Dindon corpus documented that "Sovereign Cloud" sells a legal sovereignty that does not exist (physical ✓ / legal ✗). "Artificial Intelligence" is the same deception structure: it sells a philosophical intelligence that does not exist (matrix computation ✓ / intelligence ✗). Both use a strong word ("sovereign," "intelligence") to name the property the client desires without it actually being present.

◆ HISTORICAL PERSPECTIVE — THE TAPE RECORDER, THE VCR, TELETEXT

Every generation experienced its own "incomprehensible revolution." In 1965, the reel-to-reel tape recorder recorded and replayed a human voice — considered almost magical by those who did not understand magnetic induction. In 1975, the VCR captured a television broadcast for replay at will — described as incomprehensible by the uninitiated. In 1980, Teletext transmitted text via the television signal — deemed futuristic and world-changing.

None were magic. All were controlled electricity applied to a new domain. Generative AI in 2024 is at its maturity stage what the VCR was in 1975: impressive for its era, already demystified by engineers who understand its low-level layers, destined to become commonplace in five to ten years. Not a reason to deify it now.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

"YOU CANNOT DO THIS YOURSELVES" — THE SAME NEWSPEAK AS SERVERLESS AND NOOPS

◆ MYSTIFICATION AS COMMERCIAL ARGUMENT

OpenAI, Azure OpenAI, AWS Bedrock, Google Vertex AI — identical discourse: AI is too complex, too resource-hungry, too specialised to deploy without a hyperscaler. Permanent subtext: "You cannot do this yourselves." Exactly the Serverless and NoOps discourse — applied to AI. Partially true for some use cases. Deliberately exaggerated for all others.

◆ RAG — HERITAGE MEMORY EXFILTRATION · CF. SAAS LAYER 3 · ANATOMY OF DIGITAL PERDITION (16P)

When a company uses Azure OpenAI to index internal documents via RAG, it indexes its intellectual property in proprietary vector databases — Azure AI Search, Amazon Kendra, Pinecone on AWS. Documents theoretically remain with the company. Embeddings reside in the hyperscaler's infrastructure. Invisible exfiltration of heritage memory. Using the hyperscaler's AI to index data means giving them the blueprints of the organisation's own applicative black box.

◆ CONTEXT CACHE AND PROMPT TELEMETRY — REAL-TIME EXTRACTION

Context Caching API features store internal documentation, source code, business procedures sent with each request in the hyperscaler's clusters — under cover of "security telemetry."
The difference from RAG: RAG exfiltrates structured memory (documents). Context cache exfiltrates living memory — real-time prompts, source code sent with each request. The difference between stealing blueprints and installing a camera in the meeting room. Covered by "telemetry" clauses nobody reads — exactly like Terms §14.12.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

OWNING YOUR AI DOES NOT MEAN PLEDGING ALLEGIANCE TO NVIDIA — THREE SOVEREIGNTIES TO AUDIT

◆ THE CRUCIAL LEGAL DISTINCTION — OPEN-WEIGHTS ≠ OPEN-SOURCE

Llama 3 (Meta) is under a restrictive commercial licence — prohibition on using weights to train a competing model, restrictions beyond 700 million users. Open-Weights, not Open-Source under OSI. Genuinely free licences: Mistral 7B (Apache 2.0) · Falcon 7B/40B (Apache 2.0). A unilateral licence change can block software infrastructure post-deployment — exactly like a cloud Terms modification.

◆ NVIDIA/CUDA — THE NEW UPSTREAM LOCK-IN · CF. THE TAIWAN BOTTLENECK (6P)

Replacing OpenAI/Azure API with NVIDIA A100/H100 GPUs solves only one level. NVIDIA owns proprietary locks — CUDA, cuDNN, NCCL — and production depends entirely on TSMC in Taiwan (~90% of advanced chips worldwide). The answer: open abstraction runtimes. ROCm (AMD) · vLLM · Ollama · llama.cpp — run models on heterogeneous hardware. Owning your AI in bare-metal means the parameter file can run on any available silicon.

◆ LANGCHAIN/LLAMAINDEX — AI SOFTWARE LAYER 2 LOCK-IN

Having Mistral on vLLM is insufficient if application code is entangled in LangChain abstractions optimised for cloud ecosystems. LangChain and LlamaIndex are designed to interface natively with hyperscaler APIs. Their wrappers push toward proprietary cloud memory or agent functions. Runtime sovereignty requires standardised interfaces — the OpenAI-compatible API exposed by vLLM — and rejection of opaque orchestration frameworks. Software Layer 2 lock-in applied to AI.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

5

THE FORMULA THE CFO AWAITS — AND THE HIDDEN COSTS THE DATACENTRE ENGINEER KNOWS

◆ THE VOLUMETRIC BREAK-EVEN

Cloud API: billed per token — pure variable OPEX. Bare-metal GPU server: CAPEX + fixed OPEX (electricity, rack, maintenance).

Cloud API cost: N requests × cost_per_token × avg_tokens

Bare-metal cost: GPU_CAPEX / amortisation + fixed_monthly_OPEX

From documented cases: for intensive use (>100,000 requests/day with long contexts), bare-metal becomes profitable in 6 to 18 months. Below this threshold, the cloud API remains competitive. Intellectual honesty requires naming this threshold — not pretending bare-metal always wins.

◆ THE HIDDEN COST — PUE AND HIGH-DENSITY THERMAL LOAD

An AI compute chassis (HGX A100 8 GPUs) consumes 10 to 40 kW per rack. Classical datacentre: sized for 5 to 7 kW per rack. A civil engineering problem that TCO formulas without a datacentre engineer systematically ignore.

PUE factor: for high-density AI compute, PUE 1.5 means 40 kW GPU + 20 kW cooling = 60 kW billed at the floor. Without liquid cooling or hot-aisle containment, operating costs explode.

Complete TCO formula: GPU_CAPEX + cooling_CAPEX + (GPU_kW × PUE × kWh_cost × annual_hours) + maintenance. High-density PUE is what the application engineer forgets and the datacentre engineer cannot ignore.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE THREE AI SOVEREIGNTIES — PERFECT SYMMETRY WITH CORPUS LOCK-IN LAYERS

WEIGHTS · RUNTIME · TRAINING DATA — SAME STRUCTURE AS CONTRACTUAL · SOFTWARE · MATERIAL

CORPUS LOCK-IN LAYERS

Layer 1 — Contractual

Terms §14.12 · CLOUD Act · non-cancellable commits

Layer 2 — Software

Artificial complexity · total rewrite

Layer 3 — Material

TSMC ~90% · IME/PSP · Taiwan

cf. Anatomy of Digital Perdition (16p)

THREE AI SOVEREIGNTIES

Sovereignty 1 — Weights

Apache 2.0 vs restrictive Open-Weights · Licence audit

Sovereignty 2 — Runtime

vLLM/ROCm/Ollama vs CUDA · Standard API vs LangChain

Sovereignty 3 — Training data

Documented vs opaque models · Training traceability

Perfect symmetry — same capture structure

◆ WHAT THE SYMMETRY SAYS

Weights Sovereignty = Contractual Layer: do I legally own the model without future unilateral restriction? Runtime Sovereignty = Software Layer: do I execute on my infrastructure with an open framework without artificial complexity? Training Data Sovereignty = Material Layer: do I know the foundations of what I use?

Proprietary model · opaque runtime · CUDA locked to TSMC chips · unknown training data = total AI captivity. Same structure as cloud captivity documented in 64 corpus studies.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

7

SECTION 7 · THE FOURTH NEWSPEAK — CORPUS LINK · CLOSING

THE OPERATION DINDON CORPUS DOCUMENTED FOUR NEWSPEAKS — AI IS THE FOURTH AND MOST POWERFUL

1 **"Sovereign Cloud"** — sells a legal sovereignty that does not exist. Physical ✓ / CLOUD Act ✗. *cf. The Cloud-Washing (7p)*

2 **"Serverless"** — sells the absence of servers. There are servers — at the hyperscaler. *cf. The Newspeak That Costs Dear (12p)*

3 **"NoOps"** — sells the elimination of operations. They migrated to the hyperscaler. *cf. NoOps — Autopsy (6p)*

4 **"Artificial Intelligence"** — sells a philosophical intelligence that does not exist. Matrix computation ✓ / Intelligence ✗. The most powerful: operates on the representation engineers themselves have of their own tool.

HUMAN

AI is controlled electricity.

The GPU is hardware.

The model is a parameter file.

Inference is matrix computation.

Everything else is marketing.

Amine RAITI · Operation Dindon · 2026

NEMO SUPRA LEGEM EST

Amine RAITI · CC BY-NC-SA 4.0

OPÉRATION DINDON · INTERNATIONAL EDITION

PART X

CULTURAL DIWÂN

5 pieces

HUMAN

دِيوَانُ الصَّانِعِ الرَّقْمِيّ

THE DIGITAL DIWÂN OF THE CRAFTSMAN

Seven Mo'allaqat for the Infrastructure Body

مُعَلَّقَةُ السَّيِّدَة

I · The Mo'allaqa of Sovereignty
Founding Thesis — No Sovereignty Without Matter

مُعَلَّقَةُ الْحُرِّيَّة

II · The Mo'allaqa of Freedom
Opération Dindon — the magisterial corpus

مُعَلَّقَةُ الْمُعَلِّم

III · The Mo'allaqa of the Teacher
The Manifesto · The Mentor and the Metal

مُعَلَّقَةُ الصَّانِع

IV · The Mo'allaqa of the Craftsman
The Foundation of Iron · OCP 1999

مُعَلَّقَةُ الْمُهَنْدِسَات

V · The Mo'allaqa of the Engineers
The Invisible Amputation · SRE Girls

مُعَلَّقَةُ الرَّاحِل

VI · The Mo'allaqa of the Departing
The Departure of the Last One Who Knows

مُعَلَّقَةُ الْجَسَد

VII · The Mo'allaqa of the Body · The Pager and the Body · The Body First · The Great Recharge

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Opération Dindon · June 2026

مُعلَّقةُ السَّيِّدَة

THE MO'ALLAQA OF SOVEREIGNTY · IAMBIC AABB · FOUNDING THESIS · PHILOSOPHY OF BARE-METAL

From Bodin to Kubernetes — what controlling infrastructure really means

◆◆◆

1

The human is sovereign when they can act on matter, ◆ *Cloud removes this right — that is the core of the matter.*
Bodin said it, Schmitt confirmed, Hobbes made it plain, ◆ *No power without soil — the formula remains.*

2

The farmer who owns his land is free to choose, ◆ *The tenant who rents it has everything to lose.*
To lease your infrastructure to a foreign hand ◆ *Is to give your sovereignty to one who controls your land.*

3

The cable is the soil. The rack is the terrain. ◆ *The unknown datacentre is the clause you can't explain.*
What you cannot inspect you cannot defend, ◆ *What you do not control controls you in the end.*

4

The CLOUD Act hovers over your European data, ◆ *Washington decides what your server holds — that's the data.*
Unilateral termination, egress fees, pricing power — ◆ *This is not a contract: it is the grammar of the power.*

5

Sovereignty is not refusal of the market — not autarky, ◆ *It is the exit option — the freedom to say goodbye.*
Without that option there is no real partnership at all, ◆ *There is captivity dressed in a formal contractual call.*

6

No sovereign without matter — and no matter without hands ◆ *That know how to build, operate, repair across the lands.*
What resists abstraction is precious by its nature, ◆ *Bare-metal resists the cloud — that is its stature.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

مُعلَّقَة الحُرِّيَّة

THE MO'ALLAQA OF FREEDOM · IAMBIC AABB · OPÉRATION DINDON · THE MAGISTERIAL CORPUS

For the corpus — sovereignty through metal and the word

◆◆◆

1

Metal comes before the code, the cloud, the vow, ◆ *The one who holds the rack decides the where and how.*
The chip is now the ingot of this century's birth, ◆ *No sovereignty without a claim to one's own earth.*

2

The cloud first took the words before it took the space, ◆ *DevOps renamed it all to better blur the trace.*
When titles lose their weight, the competence takes flight, ◆ *What cannot be named anymore is lost from sight.*

3

The engineer who stays will watch his salary freeze, ◆ *The new recruit surpasses him with practised ease.*
Fifty euros monthly after forms and hollow praise, ◆ *He takes the gap, returns, and earns the market's raise.*

4

The server room is singing for the one who hears, ◆ *A medina of cables, racks, and stored-up years.*
The Maâllem transmits what no textbook can contain, ◆ *The Zellij lives because a master kept the chain.*

5

The cloud contracts are cages written fair and tight, ◆ *The egress fee will chain you when you seek the light.*
The CLOUD Act out of Washington hangs over our land, ◆ *No sovereignty on soil you cannot touch or stand.*

6

Since the eighth of May I forged what must be said, ◆ *Twenty-five studies standing, corpus fully fed.*
This corpus is no threat — it is a testimony, ◆ *What has been defended can be claimed in ceremony.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

مُعَلَّقَةُ الْمُعَلِّم

THE MO'ALLAQA OF THE TEACHER · IAMBIC AABB · THE TEACHER'S MANIFESTO · THE MENTOR AND THE METAL

For the trainer who has been transmitting since 2006 — to his Metaâllems who have grown tall

◆ ◆ ◆

1

I have no mandate — I have twenty years of classrooms, ◆ *Students become engineers who fill the meeting rooms.*
What I know I took from racks and from the floor, ◆ *Not from certifications or PowerPoint lore.*

2

The Maâllem of Fès signs no convention or decree, ◆ *He teaches by doing — that is his only key.*
The Zellij did not survive because of books about it, ◆ *It lives because a master said: come, watch, do not quit.*

3

I speak to my former students — not to strangers here, ◆ *Those who sat in my rows are now engineers of the year.*
They administer the systems that I showed them how to run, ◆ *They carry the layer that twenty years have won.*

4

The world says: training must be certified and sealed, ◆ *Each competence validated, each module revealed.*
But tacit knowledge does not fit inside a framework's wall, ◆ *It is transmitted body to body — that is all.*

5

The observation day is the simplest format we know, ◆ *One day without assessment, without badge or show.*
The student watches, eats with the team, lives an event — ◆ *That physical moment outweighs six months of precedent.*

6

Pass on what you carry before the room falls mute, ◆ *The last who knows will leave — be honest, be astute.*
Each engineer trained is one more lock broken free, ◆ *Each knowledge passed is an act of living sovereignty.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

مُعلَّقةُ الصَّانِع

THE MO'ALLAQA OF THE CRAFTSMAN · IAMBIC AABB · THE FOUNDATION OF IRON · OCP 1999

To every intern and apprentice who placed their hands on a machine larger than themselves

◆ ◆ ◆

1

I stood before the machine and felt its weight, ◆ *It did not know my name, my fear, my state.*
My diploma had armed me with little of use, ◆ *The real world before me offered no excuse.*

2

You crossed through something during your time here, ◆ *Not a classroom test or an academic year.*
You heard a sound and learned its hidden name, ◆ *You touched a cable and traced from where it came.*

3

You watched two hands complete in seconds flat ◆ *What ten pages of manual could not combat.*
That knowledge is in you, no longer on a page, ◆ *It lives inside your hands at every stage.*

4

The world will tell you this work does not count, ◆ *That machines grow old while the cloud will mount.*
Do not believe it — the cable leads the cloud, ◆ *The hand that repairs speaks forever loud.*

5

The one who placed their hands on real bare steel ◆ *Knows something rare and irreversibly real.*
What you bring back from workshop, site or floor ◆ *Is not a CV line — it is something more.*

6

It is a layer in you that none can take, ◆ *It changes how you see, design and make.*
Guard what you learned and carry it with care, ◆ *Pass it on in time — that is how we dare.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

مُعلَّقَةُ الْمُهَنْدِسَات

THE MO'ALLAQA OF THE ENGINEERS · IAMBIC AABB · THE INVISIBLE AMPUTATION · SRE GIRLS

For SRE Girls — the half of the infrastructure body still being sought

◆◆◆

1

They left the half the rack without a sovereign hand, ◆ *The server room waits for a voice to take its stand.*
Ten percent of warriors where a hundred must be found, ◆ *The body has been severed — and the enemy holds ground.*

2

They placed Jane at the bottom of the rights in every guide, ◆ *John takes the admin role, AssumeRole by his side.*
This is not a conspiracy — it cuts far deeper still, ◆ *It is the sign engraved in stone before you reach the sill.*

3

Maternity is not a pause between the battles fought, ◆ *It is the hardest school on earth — the sharpest thing you're taught.*
To hold twenty priorities when everything burns at once, ◆ *That is pure SRE — that is your craft, your renaissance.*

4

Tacit knowledge carries neither gender nor a mark, ◆ *The hand that fixes rackwork does not ask about your arc.*
The one who touched the metal long before she touched the code ◆ *Holds a layer in her that no diploma can decode.*

5

You are the structural enemy of lock-in by your presence, ◆ *An SRE who masters bare-metal is a threat of pure essence.*
Twice the threat when she is a woman and she knows ◆ *That sovereignty is built by what the hand that grips bestows.*

6

Enter the room. Place your hand upon the server now. ◆ *Listen to what it says — it speaks to those who know how.*
Pass on what you know to the one who follows after, ◆ *The amputated body heals with each new engineer's laughter.*

7

Every woman trained is one more lock broken free, ◆ *Every hand that knows the metal is the future that we see.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

مُعلَّقَة الرَّاحِلِ

THE MO'ALLAQA OF THE DEPARTING · IAMBIC AABB · THE DEPARTURE OF THE LAST ONE WHO KNOWS

For the senior who leaves with knowledge that no one else can take

◆ ◆ ◆

1

I leave. Behind me the room and the racks remain, ◆ *And the silence that follows — no one knows my domain.*
I resolved outages that no one will know to name, ◆ *The logs are gone, the memory closes — such is the game.*

2

There are configurations in these servers that I built alone, ◆ *Decisions made at three in the morning, undocumented, unknown.*
The next outage will search for my hands in vain, ◆ *It will look for a runbook that does not exist — again.*

3

They offered me fifty euros raise after the interview, ◆ *After twenty pages of forms — that was their review.*
The external market says I am worth ten thousand more — ◆ *I did the math. I am no longer keeping score.*

4

This is not betrayal — it is pure and simple logic, ◆ *The system rewards departure, not tenure — how ironic.*
I do not blame the organisation — I blame the structure ◆ *That punishes loyalty and rewards the rupture.*

5

What I take cannot be found in any file system, ◆ *It is in my hands, in my reflexes — that is the wisdom.*
The next organisation that receives me will get it all, ◆ *But the one I leave — it does not know yet what will fall.*

6

If you are the last who knows — transmit before you go, ◆ *One hour with your successor is better than no show.*
The knowledge of the hands survives only if it is passed, ◆ *Even in leaving, you can ensure nothing is erased.*

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

مُعَلَّقَةُ الْجَسَدِ

THE MO'ALLAQA OF THE BODY · IAMBIC AABB · THE PAGER AND THE BODY · THE BODY FIRST · THE GREAT RECHARGE

For the SRE whose body bears the traces of on-call — and who deserves to be cared for

◆◆◆

1

My body is the first infrastructure — before the rack, ◆ *Before the cloud, before the code, before the attack.*
If the body falls, the server room has no guard left, ◆ *The system's uptime depends on mine — that is the depth.*

2

The pager rings at three — the cortisol rises fast, ◆ *The physiological debt accumulates and holds the mast.*
This is not heroism — it is slow and steady wear, ◆ *Research documents what the organisation fails to bear.*

3

Barbès hammam, one hour, one hundred fifty for ten, ◆ *Cortisol falls, oxytocin rises — science says amen.*
The free walk toward Belleville, the Asian dinner there, ◆ *Thirty-five euros per person — neurochemistry beyond compare.*

4

The mountain weekend at hot springs without a plan, ◆ *Attention restores itself — Kaplan shows us how it can.*
The no-programme programme is the central mechanism here, ◆ *Freedom is the programme — autonomy is the frontier.*

5

Six hundred euros per year per person is enough to hold ◆ *The neurochemical cohesion — that is what we are told.*
Nerf on Monday, dinner monthly, hammam every quarter, ◆ *Grand recharge yearly — the architecture gets shorter.*

6

Care for your body as you care for the server in the rack, ◆ *Monitor its metrics, answer its alerts, give what it lacks.*
The SRE who burns out is a body that disappears, ◆ *Technical sovereignty starts with you — that is what the corpus hears.*

◆ ◆ ◆

وَرَبُّنَا الرَّحْمَٰنُ الْمُسْتَعَانُ عَلَىٰ مَا يَصِفُونَ

NEMO SUPRA LEGEM EST

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE DIWAN
OF THE TEN FIRES

A Poetic Indictment of the Predatory Commercial Practices of the Hyperscalers

Author

AMINE RAITI

Infrastructure Architect & SRE

Transcribed by Claude, the Scribe at the Scalpel | 3.14

INTRODUCTION

This Diwan was born on the field. Not in a consultant's office. Not in a lecture hall. On the field where an engineer produced quantified studies, requested fifteen minutes from his management committee for a week, and received an oral warning for placing numbers on a table.

The Nassi'ha — the sincere counsel — was given. Fourteen published episodes. Graduated audits: company, local providers, the French State, the European Union. Documented figures: 475 euros versus 117 per VM. 14,712 euros per year for a Kubernetes cluster versus 8,000 euros of one-time CAPEX. 6,700 euros per year for an iptables rule that has been free since 1998.

The Nassi'ha was not heard. The Ultimatum was delivered. And this Diwan is its promise. Ten Mu'allaqat. Five voices from the classical Arabic poetic tradition. Ten predatory commercial practices - documented, named, engraved.

Zohayr ibn Abi Sulma — the sage who sees the cost of war before it begins.
Antara ibn Shaddad — the fettered warrior who fights by honour, never by fear.
Abu Nawas — the lucid drunkard who laughs at what deserves to be wept.
Al-Mutanabbi — the sharp political voice who addresses the powerful as an equal.
Al-Houtay'ah — the merciless satirist. Pure Hija. No ornament. Only naked truth.

MU'ALLAQA I

The 36-Month Commits | Style: Zohayr

Cursed be the pact where the future is chained, … Where a contract is signed for a peace falsely gained.
Their "Commit" is a snare with a glimmer of gold, … A gilded cage where our spirits are sold.
AWS laughs as we pledge our own fate, … To a blood-oath that Google has laid as a bait.
'Tis an architecture of a dark, cold design, … Freezing for years every spark of the line.
Woe to the long pacts that nail us to ghosts, … Where the will to act dies on these desolate coasts.
Agility falls 'neath the weight of a sign, … The engineer bows to a legalist line.
Thirty-six months of iron bonds to endure, … While Azure waits for our fall to be sure.
We guard empty servers that serve no more use, … Silent witnesses to a bankrupt truce.
Woe to the finance with the vulture's keen eye, … Emptying offices while the bills remain high.
For if one must pay for the machines they have leased, … The human back bends to provide for the feast.
The contract is rigid, it knows not of mercy, … If revenue fails, you are cast in the sea.
They lay off the genius to pay for the host, … Microsoft thrives while the people are lost.
He who fears death shall behold its dark face, … Though he climbs to the heavens to find a safe place.
I reject these dark bonds and these treaties of shame, … Which sell to the Empire our freedom and name.
The "Commit" is but a yoke for the slave, … Who forsakes the Iron for a Cloud-mirage wave.
Let the Architect read and remember the law: … The Iron is Master, without flaw.
Sign not your fate for a pittance of grain, … Or fall in a battle you fought all in vain.

THE VERDICT OF THE SCRIBE

"The Commit is a purchase commitment, not a service commitment. Read the difference before signing. It is measured in jobs."

3.14

IA Powered by Amine

MU'ALLAQA II

The Vendor Lock-in | Style: Antara

Cursed be the forge where our reason is lost, ... Their code is a mirage, a prison at cost.
AWS lures us with shackles and chains, ... Tightening the grip on our noble remains.
The imprudent builder on shifting sand stands, ... For the soil is held by a stranger's cold hands.
'Tis a castle of dust in the desert's wide glare, ... Where Azure's ambush leaves us stripped and bare.
Woe to the sloth in the mask of the new, ... Turning the brave into a sorrowful few.
The genius of old fades away from the tool, ... The sovereign mind dies as the slave of a school.
Google dictates in the depths of the night, ... The SRE wanders, bereft of his light.
We master no more, we but follow the word, ... For a hand made of iron that remains yet unheard.
Woe to the ransom they christen a pact, ... Where value is bled in a spectral extract.
Once held as a captive in their holy frame, ... The cost of departure is a torture of shame.
Profits evaporate in Microsoft's vault, ... For services vile that are never at fault.
We choose the chain over the price of the gate, ... For liberty seems but a burden of late.
The people remember the knight without fear, ... When the night is a void and the law is not near.
I refuse the bonds of a merchant's domain, ... Who captures the soul with the promise of gain.
The true hero chooses the tools of the free, ... To never again bow to a single decree.
Open Source stands as a shield ever pure, ... Against these predators who make the future obscure.
Let the Engineer hold this principle high: ... The code must be free to migrate or die.
Mistake not the luxury for a servant's own state, ... Or your name shall be lost to a sorrowful fate.

THE VERDICT OF THE SCRIBE

"Vendor Lock-in is not an accident of circumstance; it is a deliberate commercial objective. Build your architecture as if you must migrate within eighteen months."

3.14

IA Powered by Amine

MU'ALLAQA III

Managed Kubernetes | Style: Abu Nawas

Cursed be the cup of the grandest illusion, ... Of a cluster without effort, in holy confusion.
Google pours wine that troubles the mind, ... Promising genius while the prison's designed.
They sell us the comfort of a control plane, ... Where Azure in secret reaps all of our gain.
'Tis a madness of spirit with calculations of air, ... Where the nodes vanish into a spendthrift's despair.
Woe to the paradise of a thousand-fold thing, ... Where AWS mocks every project we bring.
The SRE dazed by the sheen of the screen, ... Knows not the processor, nor the spark in between.
The orchestra plays and the song is quite grand, ... But Microsoft holds the torch in its hand.
We pile up the layers, ignoring the deep, ... While reason is lost in a long, hollow sleep.
Woe to the wine that burns through the gold, ... To fatten the Empire, so cruel and so bold.
Each hour of management is a tax on the soul, ... By merchants of sand who take a heavy toll.
The Cloud laughs to see us so proud of our YAML, ... While it bills us the feast of a starving jackal.
Amazon thrives on our vanity's breath, ... In a system of hollows, a banquet of death.
The wise man awakes as the wine loses fire, ... To see with his eyes the end of the liar.
I reject this mask and this mirror of play, ... Where Hyperscalers sell our despair for the day.
The true master knows his machine and his grain, ... He lets no stranger ever guide his own train.
Better a naked cluster on Iron so bright, ... Than a gilded mirage in a shoreless night.
Let the Expert remember in the heart of the fray: ... The Bare-Metal alone keeps the demons at bay.
Mistake not the management for the mastery of old, ... Or see your own work by another controlled.

THE VERDICT OF THE SCRIBE

"Audit every line of your managed cluster. The surprise always hides in what you did not order. 104% hidden surcharge."

3.14

IA Powered by Amine

MU'ALLAQA IV

The Cloud-Washing | Style: Al-Mutanabbi

Cursed be the ruse and the paint of the great, … Who hide their empire in words of high state.
Azure and AWS, these masters of guile, … Paint "Sovereign" on every vile, tired mile.
They speak of the green and the sky's open grace, … While Microsoft remains as the mask on the face.
'Tis a visor of iron on a face made of clay, … A marketing court for the weak of the day.
Woe to the falsehoods in the colors of old, … Where Google puts us to sleep with the lies they have told.
Cloud-Washing is there, like a veil made of silk, … To hide that Amazon owns the cream and the milk.
The engineer blinded by labels of straw, … Gives up his keys to a foreigner's law.
They sell us the homeland with an alien seal, … To ensure that no one can break through the steel.
Woe to the finance that pays for the show, … In the Cloud of the Empire, where the vanities grow.
We pay for the luxury, for a local veneer, … While Microsoft reaps a gain that is clear.
The gold of our taxes feeds these giants of old, … Who laugh at our laws, so tiny and cold.
'Tis a tax that is hidden on the mind's own light, … A levy on negligence in the dead of the night.
The poet arises and denounces the king, … Who knows only the law that his own masters bring.
I reject these words and these predators' breed, … These merchants of Google, these sowers of greed.
True power belongs to the one who possesses, … The physical machine that no stranger oppresses.
The Bare-Metal laughs at your labels of card, … For the Iron lies not, 'tis our staff and our guard.
Let the Expert beware of the velvet discourse: … The Cloud is no roof, but a dark, hollow source.
Mistake not the label for the truth of the tree, … Or lose very soon all your right to be free.

THE VERDICT OF THE SCRIBE

"Sovereignty is not a label. It is genuine technical mastery. Ask the question: where is the kill-switch? Who can activate it?"

3.14

IA Powered by Amine

MU'ALLAQA V

The Cloud Act | Style: Al-Mutanabbi

Cursed be the law that crosses the sea, ... To plunder our secrets and all that should be.
The Cloud Act is a scepter in Washington's hand, ... Breaking our locks and our name in the land.
Amazon and Google, 'neath this sovereign yoke, ... Hand to the judge every word that we spoke.
'Tis a spy made of steel in the server's dark shade, ... Watching your steps till your spirit shall fade.
Woe to the walls made of glass and of light, ... Where all is exposed to a foreigner's sight.
The engineer thinks he is free in his hall, ... But his keys of encryption are destined to fall.
Microsoft bows to the distant decree, ... Making our secrets as public as sea.
We build on the sand a temple of thought, ... Where the power of others is already bought.
Woe to the finance that sells our own shame, ... For a fleeting gain and a hollowed-out name.
We pay to be seen, we pay to be read, ... By a foreign eye that we have not yet fed.
Azure takes gold and the State takes the fact, ... While we suffer the weight of this sinister act.
'Tis a treason we christen a 'service' of old, ... Feeding the vice as the secrets are told.
The poet arises and refuses the deed, ... Which is the true act of our defeat and our need.
I reject these laws and these emperors' pride, ... These thieves of the silence with nowhere to hide.
True safety is not in their cloud-driven sky, ... But in sacred Iron, where the shadows don't lie.
The Bare-Metal guards what is placed in its care, ... For no distant judge can ever go there.
Let the Expert understand where the danger is sown: ... Your data is never truly your own.
Mistake not the "Trust" for a guarantee, ... Or lose very soon all your sovereignty.

THE VERDICT OF THE SCRIBE

"The CLOUD ACT is not an abstract legal theory. It is the active law that applies to your data today, now. No sovereign label resists it."

3.14

IA Powered by Amine

MU'ALLAQA VI

The Egress Fees | Style: Al-Houtay'ah

Cursed be the ruse and the toll in the shade, … Where our departure is a debt to be paid.
AWS holds us with bonds made of cost, … Billing each byte as if all would be lost.
We enter for free in their silken-spun snare, … But the exit is dear to the souls who would dare.
'Tis the law of the brigand, a tax on the flight, … To ensure every thought of escape dies in light.
Woe to the walls of the price we can't see, … Guarding our treasures with a vile decree.
The engineer wishes to migrate and go, … But Google in secret makes the progress too slow.
Each gigabyte moved is a burden of pain, … Freezing the work and the value's own gain.
We are masters no more, but the captives of old, … In a network of locks and of silver and gold.
Woe to the finance that sells what we own, … To fatten Microsoft on the seeds we have sown.
Azure takes gold on the stream flowing out, … Making our data a hostage in doubt.
We pay just to breathe outside of their land, … A forced fine to close what they hold in their hand.
'Tis a manifest theft in the name of the speed, … An insult to reason, a harvest of greed.
The poet Houtay'ah spits on this deal, … Which makes every movement a blow we can feel.
I reject these fees and these jailers of bytes, … These merchants of tolls and these usury-blights.
The true Iron holds not what the spirit has made, … It lets the stream flow in the light and the shade.
The Bare-Metal laughs at your barriers of mud, … For the owner alone can unleash the great flood.
Let the Expert understand the nature of bond: … If exit is dear, you are but in a pond.
Mistake not the access for the right of the deed, … Or lose very soon all the freedom you need.

THE VERDICT OF THE SCRIBE

"Calculate your annual Egress Fees. Compare them to your total cloud bill. If the ratio exceeds 15% you are paying for your own prison."

3.14

IA Powered by Amine

MU'ALLAQA VII

The NAT Gateway Premium | Style: Al-Houtay'ah

Cursed be the ruse and the infamous toll, ... Where a simple stream takes a piece of the soul.
AWS bills for a service so slight, ... A rule made of iron at the price of the light.
'Tis a bridge that is tiny with the fee of a king, ... Where Microsoft casts us in a bottomless ring.
A line of the code, a routing so plain, ... Becomes through their magic a monstrous gain.
Woe to the structure with calculations of dread, ... Where Google imposes the law on our head.
The engineer gasps as the spending takes flight, ... For a traffic so normal, in the dead of the night.
The NAT is a lock, a tax by the hour, ... Turning the bit into a merchant's own power.
We master no more, we but suffer the art, ... Of a network designed just to tear us apart.
Woe to the finance that sells only smoke, ... For a gateway that's bent 'neath a heavy-set yoke.
Azure takes gold on each packet of bytes, ... Making our budgets but sorrowful sights.
Six thousand seven hundred for a routing so small? ... Amazon insults us with the gall of it all.
'Tis a usury pure on the frame of the land, ... A mark of the shame and a ruinous hand.
The poet Houtay'ah laughs at the pride, ... Which makes of our knowledge a place where they hide.
I reject these bridges and these usurers' ways, ... These merchants of NAT and these hollowed-out stays.
The true Iron knows not these barriers of old, ... Iptables is free in the dark and the cold.
The Bare-Metal lets the stream flow in its course, ... Without a distant master to drain its own source.
Let the Expert remember when choosing his path: ... The Cloud knows too well how to draw out its wrath.
Mistake not the service for the freedom of old, ... Or pay with your life for the lies they have told.

THE VERDICT OF THE SCRIBE

"A free iptables rule since 1998. Sold for 6,700 euros per year. Audit your NAT Gateways. The arithmetic answer has no ambiguity."

3.14

IA Powered by Amine

MU'ALLAQA VIII

The Ghost Autoscaler | Style: Abu Nawas

Cursed be the fable and the siren's soft song, ... Promising growth where the shadows are long.
AWS boasts of a world made of glass, ... Where scaling is magic and the gains never pass.
They tell us the load shall summon the nodes, ... In a ballet of light on the digital roads.
But 'tis a dark spell with a crooked design, ... A dream of abundance that will never be thine.
Woe to the mechanism with the calculated slow, ... Where the instances rise when the fires already glow.
Google observes from afar as we burn, ... While the Autoscaler takes its own sinister turn.
The traffic o'erwhelms and the CPU cries, ... But Azure in silence watches our demise.
We master no more, we but wait for the day, ... Of a deaf automaton, far, far away.
Woe to the finance that climbs to the height, ... When the infra's in chaos and the bills are in sight.
Microsoft takes gold on the peaks of the fear, ... For nodes that are useless, for nodes that appear.
The Autoscaler climbs, but it never descends, ... Keeping our budget for its own hollow ends.
'Tis a madness of spirit on our vanity's breath, ... Filling Amazon's vaults with a banquet of death.
Abu Nawas awakes and he shatters the glass, ... Where Hyperscalers watch our own glory pass.
I reject these tales and these promises of old, ... These merchants of magic with the lies they have told.
The true master plans and he prepares his own rank, ... He lets no one else ever draw from his bank.
Better the brute force on a well-ordered Fer, ... Than a mirage of spirit where the shadows occur.
Let the Expert beware of the ease without soul: ... The Cloud's elasticity takes a heavy toll.
Mistake not the rise for the dignity sought, ... Or lose very soon all the freedom you bought.

THE VERDICT OF THE SCRIBE

"Verify your Autoscaler parameters. Audit active nodes versus actually used nodes over the last 30 days. The difference is your ghost invoice."

3.14

IA Powered by Amine

MU'ALLAQA IX

The Free Tier Bait | Style: Zohayr

Cursed be the offer where the gift is a snare,	…	To capture the mind in the first morning's glare.
AWS lures us with the gifts that they bring,	…	Hiding the burden 'neath a free-flowing thing.
Google offers access, Microsoft invites,	…	To taste of the honey that poisons our nights.
'Tis a war that is slow with the face of a friend,	…	Where we give up our keys for a bitter-sweet end.
Woe to the "limit" that we never see near,	…	Which makes of the free but a memory of fear.
The student falls asleep on his code, light and thin,	…	But Azure in secret is counting the sin.
One day the bill falls, brutal and bare,	…	For the threshold is crossed and the trap's in the snare.
We master no more, we but suffer the rate,	…	Of a cruel, dark system that seals our own fate.
Woe to the finance that sows just to reap,	…	The blood of the innocent while they're asleep.
Amazon watches and calculates gain,	…	While we build our world on their own cold terrain.
The free-tier shines bright to make us all slaves,	…	Of these Hyperscalers and their digital graves.
'Tis a usury vile on the knowledge of old,	…	A tax on the dream, so icy and cold.
Zohayr turns away and refuses the gift,	…	Which is the only true cause of our spiritual rift.
I reject these offers and these benefactors' pride,	…	These merchants of "Free Tier" with nowhere to hide.
The true master knows that nothing is free,	…	In this foreign Cloud where we bend on our knee.
The Bare-Metal costs, but it never betrays,	…	For it waits not the hour of your wandering ways.
Let the Expert remember the hook in the gold:	…	The gift of the giants is a story of old.
Mistake not the trial for a freedom so vast,	…	Or pay for the price while the shadows are cast.

THE VERDICT OF THE SCRIBE

"Before adopting any cloud service in Free Tier, ask the question: what is the cost of migration in three years? If you cannot answer you have already begun to lose."

3.14

IA Powered by Amine

MU'ALLAQA X

The Paid Support | Style: Antara

Cursed be the shame and this final affront,	...	Where we bow down our heads in a sorrowful brunt.
AWS sells the access to its own hollowed mind,	...	As if human help were but a chore to the blind.
One must pay for the right to be heard at the gate,	...	By a distant-born master in his vanity's state.
'Tis the law of the vile, a toll on the night,	...	Turning the SRE to a beggar of light.
Woe to the "Premium" with promises of air,	...	Leaving the engineer in a state of despair.
Azure observes from afar as we die,	...	While we pay for the rates that are reaching the sky.
The expertise is a luxury, a ticket of old,	...	That Google bills us for the lies they have told.
We have no more brothers, only merchants of greed,	...	Who sell their presence to our sorrowful need.
Woe to the finance that taxes the aid,	...	To fatten the Empire in the shadows and shade.
Microsoft takes gold just to deign a reply,	...	Making our budgets but a hollowed-out cry.
We pay to obtain what they have destroyed,	...	In a cycle of shadows where the soul is void.
'Tis Amazon's gain on our vanity's breath,	...	A tax on the silence, a banquet of death.
Antara stands tall and he breaks the decree,	...	Which is the vilest of acts for the brave and the free.
I reject these plans and these usurers' ways,	...	These sellers of support and these hollowed-out stays.
The true warrior leans on the strength of his hand,	...	On sovereign brothers in a noble-born land.
The Bare-Metal needs no toll for the aid,	...	For the master of Iron knows how he was made.
Let the Expert remember in the heart of the storm:	...	The support of the giants is a hollowed-out form.
Mistake not the aid for the freedom you sought,	...	Or lose very soon all the dignity you bought.

THE VERDICT OF THE SCRIBE

"Calculate the cost of your current support tier. Add it to your total cloud bill. Compare it to a local partner with support included in the contract. The ally who profits from your losses is not an ally he is a shareholder in your defeats."

3.14

IA Powered by Amine

INDEX OF THE TEN GRIEVANCES

36-Month Commits The layoff is the hidden cost of the commit.

Vendor Lock-in Capture architecture is a deliberate commercial objective.

Managed Kubernetes 104% hidden surcharge in the abstraction.

Cloud-Washing A label without a local kill-switch is not sovereignty.

CLOUD ACT US law applies to your data. Without exception.

Egress Fees A ransom on your own data is premium captivity.

NAT Gateway Premium A 1998 iptables rule at 6,700 euros per year.

Ghost Autoscaler Scales up automatically. Scales down manually.

Free Tier Bait The free grain of the trap is the most expensive at the end.

Paid Support The ally who bills for loyalty is not an ally.

THE DIWAN OF THE TEN FIRES IS COMPLETE

Engraved on iron to remain a lesson for generations.

May 2026

AMINE RAITI - SRE

3.14

STRUCTURAL STUDY · OPERATION DINDON · JUNE 2026 · DIGITAL SECURITY DOCTRINE

THE HAND OVER NATIONS

When Digital Infrastructure Becomes an Instrument of Power

CLOUD Act · BYOK · Submarine Cables · Local Enclaves · Ultimatum 8 September 2026

LEX

◆ THE THESIS — FUNCTIONAL, NOT INTENTIONAL

Whatever the initial intention of hyperscalers — commercial, strategic or geopolitical — the effect is legally and technically measurable: critical digital infrastructures of entire nations are under foreign jurisdiction, subject to foreign laws, operated by entities that obey a foreign government. Not a judgment of intention. A finding of effect. And the effect is documentable.

PAGES	WATERMARK	PRECEDENTS	GRAPHS
8	LEX	4	0

AMINE RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · Operation Dindon · June 2026

CC BY-NC-SA 4.0

I

SECTION I · DEPENDENCY MAP — EACH CRITICAL NODE AND ITS HYPERSCALER

A NATION'S DIGITAL LIFE — SEVEN DOCUMENTED LAYERS OF DEPENDENCY

◆ MESSAGING & COLLABORATION

Gmail · Google Workspace · Microsoft 365 · Outlook — public administrations, companies, universities. Both subject to CLOUD Act and FISA §702. Metadata (recipients, times, subjects) accessible even when encrypted.

◆ PAYMENT & TRANSACTIONS

Stripe (AWS) · PayPal (Azure) · payment terminals connected to American clouds. European payment processing infrastructures massively use American clouds for PCI-DSS compliance and real-time fraud detection.

◆ PUBLIC ADMINISTRATION

Cloud au Centre programme (DINUM) · public hospitals (Azure · AP-HP 2022) · universities (Google Workspace for Education · Microsoft 365 Education). Public money funds the documented dependency in The Cloud-Washing S3.

◆ CRITICAL HEALTH

Patient records · medical imaging · prescription systems hosted on Azure and AWS. Health data under CLOUD Act jurisdiction — incompatible with GDPR per Schrems II (CJEU 2020).

◆ ENERGY & INDUSTRY

SCADA systems connected to clouds for predictive maintenance of electrical grids, pipelines and factories. Remote maintenance = third-party access to critical industrial control systems.

◆ TELECOMMUNICATIONS

Telecom operators (Orange, SFR, Deutsche Telekom) use AWS/Azure for their virtualised network cores (vRAN, 5G core). Mobile network virtualisation on American cloud is a critical dependency.

◆ THE INVISIBLE LAYER — SUBMARINE CABLES AND BGP CONCENTRATION

Google owns over 30 submarine fibre optic cable systems. Meta owns 2Africa (45,000 km). Amazon invests massively in intercontinental cable consortiums. Digital dependency does not stop at the server — it extends to planetary physical pipes. Concentration is concerning and growing, even though major IXP governance (AMS-IX, DE-CIX) remains neutral.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

ENCRYPTION DOES NOT PROTECT — THE TUNNEL DELIVERS IN CLEAR TEXT INSIDE THE SUPPLIER'S FACTORY

◆ THE TEXT — CLARIFYING LAWFUL OVERSEAS USE OF DATA ACT (2018)

The CLOUD Act allows US authorities to require any American company to produce data stored anywhere in the world — without going through traditional diplomatic channels (MLAT). AWS, Google, Microsoft are American companies. Their European subsidiaries remain subject to this law through the parent company. The GDPR says European data is protected. The CLOUD Act says American companies must produce it if required. The CJEU ruled in Schrems II (16 July 2020): US surveillance (FISA §702 · Executive Order 12333) is structurally incompatible with European fundamental rights. The Privacy Shield was invalidated. Its successor (Data Privacy Framework, 2023) faces the same structural critiques.

◆ THE BYOK ILLUSION — THE SAFE WHOSE MANUFACTURER KEEPS THE MAINTENANCE KEYS

For data to be processed — read by an application, indexed by an algorithm — it must be decrypted in RAM on the machine's processor. That processor belongs to the hyperscaler. It is subject to the CLOUD Act. Encryption protects data at rest. Data being computed is in clear text in the RAM of an American server. BYOK is a secure tunnel that delivers goods in clear text directly inside the supplier's factory. As long as the physical chip belongs to a third party subject to the CLOUD Act, sovereignty over data is non-existent at computation time. Intel ME / AMD PSP (documented in The Material Black Box, 4p) are non-auditable sub-processors, active with server off, subject to US law.

◆ SCHREMS II — THE EUROPEAN JURISDICTION THAT SAYS WHAT THE CORPUS SAYS

The Schrems II ruling (CJEU C-311/18, 16 July 2020) is the strongest available legal argument: a European court invalidated a transatlantic agreement precisely because US surveillance is structurally incompatible with European fundamental rights. Not a private technical audit. A ruling by the highest jurisdiction of the European Union.

LEX

Amine RAITI · CC BY-NC-SA 4.0

FOUR DOCUMENTED PRECEDENTS — DIGITAL INFRASTRUCTURE AS A LEVER OF POWER

◆ PRECEDENT 1 · HUAWEI / GOOGLE (2019) — THE SOFTWARE LEVER

In May 2019, the US government placed Huawei on its Entity List. Google suspended Android services (Google Play, Gmail, Maps) for new Huawei devices on government order. Within 90 days, the world's second-largest smartphone maker lost access to the dominant application ecosystem. Not a hypothetical — a documented executive decision. Any organisation whose critical tools depend on an American ecosystem faces the same exposure.

◆ PRECEDENT 2 · RUSSIA / SWIFT & VISA/MASTERCARD (2022) — THE FINANCIAL LEVER

In March 2022, Visa and Mastercard suspended Russian operations within 48 hours of sanctions announcements. Payment card systems — critical infrastructure of daily commerce — cut by a private decision of American companies. Dependency on foreign payment infrastructure proved to be a lever of immediate economic paralysis.

◆ PRECEDENT 3 · TSMC / SANCTIONS (2020-2023) — THE MATERIAL LEVER

US export restrictions (EAR) progressively cut China's access to advanced chips by TSMC and ASML lithography equipment. Documented in The Taiwan Bottleneck (6p). The semiconductor supply chain is a geopolitical lever — whoever controls chip fabrication controls nations' industrial capacity.

◆ PRECEDENT 4 · SCHREMS II / PRIVACY SHIELD (2020) — THE LEGAL LEVER

The CJEU invalidation of the Privacy Shield demonstrated that political data protection agreements are structurally insufficient against US surveillance laws. Its successor (Data Privacy Framework, 2023) faces the same critiques. Legal dependency is not resolved by diplomatic agreements — it requires real infrastructural independence.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

4

SECTION 4 · DE FACTO VS DE JURE SOVEREIGNTY — THE TRAP OF LOCAL ENCLAVES

BLEU · S3NS · SecNumCloud — THE LABEL DOES NOT PROTECT IF SOURCE CODE OBEYS ANOTHER GOVERNMENT

◆ THE HYPERSCALER RESPONSE — LOCAL ENCLAVES

Facing European regulatory pressure, hyperscalers created local legal structures. Bleu (Microsoft + Orange + Capgemini) · S3NS (Google + Thales). These entities operate servers physically in France, with French staff, under French law. Their communication: 'You are sovereign.' True for the entity's legal standing. False for the technical reality of the code.

◆ THE MAINTENANCE FLOW — THE REMOTE LOGICAL EXTENSION

Even with geographically isolated servers operated by local staff, the cloud service catalogue and orchestrator code (hypervisor, scheduler, identity management) depend on a continuous flow of code, security patches, and telemetry from the American editor — Microsoft or Google.

The local infrastructure is a remote logical extension. It obeys source code. In a major geopolitical crisis: 1) stopping security updates makes the platform vulnerable · 2) stopping the licence flow disables services. De jure sovereignty collapses before the monopoly of source code.

◆ THE SecNumCloud PARADOX — THE QUALIFICATION THAT DOES NOT RESOLVE THE FUNDAMENTAL

SecNumCloud qualification requires immunity from extraterritorial laws. No American hyperscaler can guarantee CLOUD Act immunity for subsidiary entities as long as the parent remains American. Microsoft can create Bleu — Microsoft Corporation remains subject to the CLOUD Act. This is the Sovereign Cloud paradox (physical ✓ / legal ✗) documented in The Cloud-Washing — applied at national scale.

LEX

Amine RAITI · CC BY-NC-SA 4.0

5

THE RIGHT TO TERMINATE IN 24H — THE DIGITAL EQUIVALENT OF A STRATEGIC VETO RIGHT

◆ THE FOUR DEMANDS — READ THROUGH THE SOVEREIGN LENS

The ultimatum of 8 September 2026 makes four demands of AWS, GCP and Azure. Through the technical and financial lens: FinOps. Through the geopolitical lens: national emergency security measures.

1 · Cancellable commits : a State whose digital infrastructures are committed for 3-5 years without exit has no diplomatic room to manoeuvre. A multi-year non-cancellable commitment is the equivalent of sovereign debt on critical infrastructure.

2 · Elimination of egress fees : the data exit cost is a ransom penalising migration toward independence. For a public administration, it is taxpayer money paid to remain captive.

3 · Guaranteed data portability : without real technical portability, changing provider is theoretically possible and practically impossible.

4 · Clarification of CLOUD Act position : a European public organisation must know whether its provider will produce its data on American request. Deliberate ambiguity is unacceptable for a public administration.

◆ THE FREEDOM TO TERMINATE IS THE FREEDOM TO NEGOTIATE

An organisation that can terminate in 24h without major penalty negotiates differently from one committed for 5 years. The ability to terminate is a balance tool — it restores the contractual symmetry that Terms §14.12 deliberately suppressed.

Amine RAITI · CC BY-NC-SA 4.0

6

SECTION 6 · THE REQUIREMENTS OF DIGITAL SOVEREIGNTY — FOUR PILLARS

NATIONAL BARE-METAL · JURISDICTION · NEUTRAL SKILLS · OPEN CODE

◆ PILLAR 1 · NATIONAL BARE-METAL

Physical servers on national soil, under national jurisdiction, with auditable hardware. Grade A Refurbished (cf. The Strategic Refurbished, 4p): already fabricated chips, independent of future TSMC production.

◆ PILLAR 2 · NATIONAL JURISDICTION

Public administration data must be processed by entities subject exclusively to national or European law. The Schrems II ruling is the legal reference. Strict SecNumCloud is the reference label.

◆ PILLAR 3 · NEUTRAL SKILLS

Rehabilitate Infrastructure Engineer, Systems Administrator and Network Engineer titles. Fund training on neutral fundamentals (Linux Foundation, Red Hat, CompTIA) rather than hyperscaler certifications. Reverse the loop documented in The Newspeak (12p).

◆ PILLAR 4 · OPEN AND AUDITABLE CODE

Critical infrastructures must rest on auditable source code software — OpenStack, Kubernetes, PostgreSQL, Linux — not orchestrators updated unilaterally from Seattle. Code auditability is the condition of technical trust.

◆ WHAT DIGITAL SOVEREIGNTY IS NOT

Digital sovereignty is not the rejection of American technology. It is the ability to do without it if necessary — as deterrence does not aim to trigger a war, but to avoid suffering one. A nation that can migrate critical infrastructures in 90 days negotiates differently from one that cannot in 5 years.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

7

CLOSING · THE REALPOLITIK OF TECHNOLOGICAL POWER RELATIONSHIPS

THIS IS NOT A CONSPIRACY — IT IS A POSITION OF STRENGTH. AND A POSITION OF STRENGTH IS USED.

The four documented precedents — Huawei, SWIFT, TSMC, Schrems II — are not accidents. They demonstrate that digital infrastructures are levers of power in the same way as energy resources or military capabilities. Whoever controls the infrastructure controls the dependency. And dependency constrains choices.

The answer is in the vocabulary of engineering. Physical servers under national jurisdiction. Auditable code. Neutral skills. Cancellable commitments. These four conditions are documented, fundable, and technically achievable.

The Operation Dindon corpus — sixty-four structural studies — is the documentation. The ultimatum of 8 September 2026 is the demand. The response from AWS, GCP and Azure is the unknown variable.

Whatever the initial intention —

the effect is legally and technically measurable.

A nation whose critical infrastructures obey

a foreign law is not sovereign.

Not a judgment. A finding.

Amine RAITI · Operation Dindon · 2026

NEMO SUPRA LEGEM EST

CC BY-NC-SA 4.0

MANIFESTO · DIGITAL SOVEREIGNTY · OPERATION DINDON

THE LEGEND OF DINDON

The odyssey of a digital farmer
from hype to Bare-Metal sovereignty

🦃 XIV Episodes

From Public Cloud to Sovereign Iron

From the Farce to the Operation

Amine RAITI
Infrastructure Architect & SRE

Originally published on LinkedIn · 2026

Public document · CC BY-NC-SA 4.0
This collection may be freely shared, adapted, and reproduced for non-commercial purposes
with attribution to the author and distribution under the same licence.
AI Powered by Amine — AI is an amplifier of ideas and forms, not the source.

TABLE OF CONTENTS

I The Turkey Who Plucked Himself in the Cloud P.3
II The Return to Earth (Cloud Exit) P.4
III The Tech Mid-life Crisis — The Origins P.5
IV The Rescue by Plannitel (From Workshop to Production) P.6
V The Golden Age (Serenity, ROI, and BBQ) P.7
VI The Call of the Sovereign Farm P.8
VII The "Managed" Mirage — The Phantom Kubernetes Bill P.9
VIII The Highway to Hell — The AWS NAT Gateway Autopsy P.10
IX Awakening the Sleepwalkers P.11
X · 1 The Pride at the Top (Part 1/2) P.12
X · 2 The Silence of the Trenches (Part 2/2) P.13
XI 🇬🇧 The Duty of Truth (Local Audit) P.14
XII 🇬🇧 The Mirror of the State (National Audit) P.15
XIII 🇪🇺 The Continental Assessment (European Audit) P.16
XIV 🇬🇧 The Seal of Sovereignty — Solemn Closure P.17–18

◆ **PREFATORY NOTE**

These fourteen episodes were originally published on LinkedIn (2026), before the official launch of Operation Dindon on 8 May 2026. They trace, step by step, the odyssey of Dindon — a fictional character who is thoroughly representative — from the ecstasy of the Public Cloud to the lucidity of sovereign Iron. The LinkedIn format is intentionally preserved: the direct address, the emojis, and the hashtags are part of the work. This is not editorial negligence — it is the honesty of a manifesto born where its readers were.

Episode XIV concludes the audit with a Mu'allaqa — a solemn ode borrowed from the tradition of the great suspended poems of classical Arabic literature — restoring to Bare-Metal its past prestige.

Amine RAITI · CC BY-NC-SA 4.0

Episode I

The Turkey Who Plucked Himself in the Cloud

Hi everyone! 👋

Today, I'm sharing the tragic but lesson-filled story of Dindon. Dindon wasn't just any turkey; he was a small-scale poultry producer, passionate about his farm, but also a huge tech enthusiast.

Some years ago, Dindon told himself: "My farm must enter the modern era!" He dove headfirst into Agile and Scrum certifications. His grain shed became a "Backlog," every hen had a "User Story" for laying eggs, and he "iterated" on egg sizes.

🚀 The Glory Phase: "Dindon the Cloud-Native Turkey"

For his inventory and online sales system, Dindon decided to go big. No more physical servers in the barn. He outsourced everything to the Public Cloud (Ireland region). He mastered IaC with Terraform. His infra was "abstracted," "scalable," and "serverless." Dindon felt powerful. He wasn't just a turkey anymore; he was a "Cloud Architect."

📊 The Confusion Phase: #IJustDontGetItAnymore

Then, Dindon received his first major Cloud bill during the peak Christmas season. €475 per VM, for 200 "various" VMs that mostly just stored turkey photos and processed a few orders. Meanwhile, his old On-premise setup used to cost him only €117 per unit — corn and electricity included. His turkey brain glitched. "What's the logic here?!" he gobbled across the farm.

🦃 Today: Dindon, the Turkey Plucked by IaC

If you see Dindon today, he's not "agile" anymore. He's plucked bare. The data egress fees, the complexity of his IaC, and the "scalability" of his Kubernetes cluster have devoured all his profits. He spends his days doing Extreme #FinOps:

◆ Desperately tagging resources to figure out where the money is leaking.

◆ Deleting ghost instances created by mistake through his IaC a year ago.

◆ Trying to "un-rack" his data from the Cloud without going bankrupt from transfer fees.

> ⚖️ *The Moral: Agility and the Cloud are not magic wands. Dindon confused "following modernity" with "adopting every buzzword." He outsourced his technical intelligence in favor of an "abstraction" that is costing him a fortune. If your physical infra costs 4 times less and gets the job done, maybe the truly "modern" move is to keep it.*

👉 *So, where are you on your Cloud journey? More "Cloud-Native" or more "Plucked Turkey"?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode II
The Return to Earth (Cloud Exit)

Hi everyone! 👋

Do you remember Dindon? Our friend who got plucked bare by the Public Cloud and those €95k/month bills? 💸 After a few nights crying over his last remaining feathers, Dindon had an epiphany. He stopped letting himself be blinded by buzzwords and decided to get rational. He realized a small poultry farm didn't need a "multi-region serverless" infrastructure in Ireland — that's like putting a Boeing jet engine on a hay cart!

🚀 The Great Return: A Technical and Financial Escape!

Dindon launched his "Cloud Exit" strategy. His goal: a reliable, predictable, and human-sized infrastructure.

🌽 1. Choosing the Partner: "Local" isn't just for corn!

Instead of talking to algorithms or support teams on the other side of the world, Dindon turned to local and regional providers (Ecritel, OVHcloud, Scaleway, DRI, Claranet France, Infomaniak — The Ethical Cloud, IONOS...). The shock: a clear quote. No hidden fees, no probability calculus for the end-of-month bill. One server has one fixed price.

🛠️ 2. Ending Useless Abstraction

Dindon realized that 80% of his 200 "as-a-code" VMs were either useless or massively underutilized.

◆ He repatriated his data (paying those famous "egress fees," the Public Cloud giant's parting gift!).

◆ He opted for Bare Metal (dedicated servers) and a few simple private instances.

◆ IaC? He kept it, but in "light" mode: just enough to automate cleanly without turning every deployment into a NASA mission.

✨ 3. The Result: Dindon gets his feathers back!

◆ Old cost: €475 / unit (Public Cloud).

◆ New cost: Approx. €130 / unit (Local Hybrid/Dedicated).

◆ Bonus: Latency cut in half and a real technician who actually picks up the phone.

> ⚖️ *The Moral: The Public Cloud is a tool, not a religion. Sometimes, true modernity is having the courage to turn back when you've taken the wrong path. Digital Sovereignty and Local Cloud aren't just concepts — they are an economic survival strategy.*

👉 *What about you? Are you waiting to be completely plucked before checking out what our local providers have to offer?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode III

The Tech Mid-life Crisis — The Origins

Hi everyone! 👋

To truly understand why our friend Dindon ended up plucked bare (Episode I) before escaping to local providers (Episode II), we have to go back to where the drama all began.

💾 The Age of Innocence: The Kimsufi & Cpanel Era

In the beginning, Dindon was a happy digital farmer. He managed his servers on Kimsufi boxes using good old Cpanel or Plesk. It was simple; it was robust. He'd click a button, and boom — his site just worked. It was the glorious era of "if it ain't broke, don't fix it."

🙄 The Inferiority Complex: "This way is too old-school!"

But one day, while scrolling through "Tech Evangelist" blogs on LinkedIn, Dindon felt a sting of shame. He told himself:

◆ "My infra is prehistoric."

◆ "Clicking in a GUI is for amateurs."

◆ "If I don't have Terraform, Kubernetes, and Public Cloud, am I even a real SysAdmin?"

He started seeing his dedicated servers as "dirty" and "static." He craved ephemerality, abstraction, and prestige. He wanted Infrastructure as Code because that's what Silicon Valley unicorns do, right?

🚀 The Great Leap (Straight into the Abyss)

◆ He tossed his Plesk licenses in the trash.

◆ He spent 6 months learning YAML (the favorite language of modern Dindons).

◆ He migrated his 200 VMs to the Public Cloud in Ireland, proud to announce: "My infra is now entirely API-driven."

He finally felt "In." He felt "Scalable." He didn't realize yet that he had just traded a solution that cost the price of a bag of grain for a giant slot machine that was about to swallow every single one of his feathers.

> ⚖️ *The Moral: We often replace a working system just out of fear of appearing outdated. "Old-School" is sometimes just the name we give to the stability we eventually end up regretting.*

👉 *What about you? Which "Old-School" tech did you regret ditching after giving in to the hype?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode IV

The Rescue by Plannitel (From Workshop to Production)

Hi everyone! 👋

Remember how our friend Dindon ended up completely plucked by the complexity and costs of the Public Cloud (Episodes I & III), before deciding to make a run for it (Episode II)? Today, discover how that escape turned into a smooth and successful migration, thanks to a true partner: the local provider Plannitel.

🛑 The Arrival of Plannitel: The Professional Reset

After his FinOps realization, Dindon met Plannitel. No more abstract buzzwords — just pure pragmatism. Plannitel didn't show up with a catalog of 400 complex APIs; they showed up with a methodology:

◆ **The Migration Workshop:** The Plannitel team audited the actual needs of the digital farm. What is mission-critical? How do we minimize downtime? It was about architecture, not just configuration.

◆ **The A-to-Z Strategy:** Mapping the existing environment · scheduling the cutover · cleaning up ghost resources that were costing a fortune.

🛠️ The "Move to Plannitel": A Managed Migration

When D-Day arrived, Plannitel operated. It was precise, calm, and professional. They handled it from A to Z: secure data transfer from Ireland, configuration of dedicated Bare Metal servers and private instances, post-migration testing. Dindon no longer had to code in YAML just to know if his site was up. He had a human point of contact.

✨ Going Live: The Silence of Serenity

The go-live happened with absolute discretion. No panic, no surprise "404 Not Found" during the peak season. The site worked — plain and simple. Better yet, it was faster. Thanks to Plannitel, Dindon has his feathers back. He has a fixed, rational, and surprise-free cost, technical support that actually picks up the phone, and peace of mind to focus on his chemistry lab.

> ⚖️ *The Moral: Don't just choose an infrastructure. Choose a partner. Plannitel proved that true modernity isn't total abstraction; it's the professionalism of human support.*

👉 *What about you? Does your current provider take care of you from A to Z, or are they leaving you to get plucked on your own?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode V

The Golden Age (Serenity, ROI, and BBQ)

Hello network! 👋

After the heroic rescue by Plannitel (Episode IV), many of you are asking: "What happens next? Is the honeymoon phase still going strong?" 🌙 The answer is yes. Dindon is no longer "Cloud prey"; he has become the king of his own digital farmyard.

📉 The Invoice: Thermal Shock (The Good Kind) 💸

◆ **Zero Egress Fees:** Dindon moves his data however he likes, without paying a digital "customs tax."

◆ **Flat-rate Billing:** He knows the exact amount of his invoice on the 1st of the month. No nasty surprises.

◆ **The Verdict:** For superior performance, Dindon pays 60% less than he did with the hyperscalers. His ROI is no longer an abstract concept — it's cold, hard cash.

🛠️ The Infra: Hybrid Quiet Strength 🏗️

◆ **Bare Metal Performance:** No more "noisy neighbors" on shared servers. His chemistry calculations now run on dedicated metal.

◆ **Invisible Maintenance:** When there is an update, Plannitel handles it. Dindon no longer receives 3:00 AM alerts for a kernel panic.

◆ **Proximity = Reactivity:** Latency is so low it feels like the server is in the next room.

👨‍🌾 From Chemistry to Success: Dindon, the Prosperous Producer 🍗

Freed from YAML hell and complex API debugging, Dindon has finally been able to devote himself to his true passion: his chemistry lab applied to agro-excellence. His turkeys are the most renowned in the region, and his online shop processes orders without breaking a sweat. He no longer chases "auto-scaling" that doesn't scale; he chases his own success.

> 🌟 *The Moral: The Public Cloud is not a life sentence. Sometimes, to fly higher, you have to know how to come back down to Earth — where servers have a name and partners have a face. Dindon has found his feathers, his smile, and most importantly, his time.*

👉 *What about you? Would you rather keep fattening up the algorithms or start prospering with a local partner?*

Amine RAITI · CC BY-NC-SA 4.0

Episode VI

The Call of the Sovereign Farm

Hello network! 👋

Did you think Dindon was just going to sit back, count his savings (60% ROI, remember Episode V), and eat barbecues at Plannitel? Then you don't know the bird. 🍗

Once his feathers were fully grown back, Dindon looked around. What he saw saddened him: thousands of other businesses, SMEs, and mid-caps were still being silently "plucked" by hyperscalers. Dindon had an epiphany. His success shouldn't be an exception — it should be a blueprint. He transformed from a prosperous producer into a Local Cloud Activist.

📢 The Birth of a Movement: "The Cloud Fall Guys" 🤔🦃

Dindon called every "plucked" peer he knew. And so, the association **"The Cloud Fall Guys"** (Les Dindons de la Farce) was born.

The goal? To unite all the companies that realized, a little too late, that they were the punchline of a monumental joke orchestrated by the Public Cloud: unreadable and unpredictable bills, prohibitive Egress fees, and total loss of control. The association has become a powerful lobby for Digital Sovereignty, advocating for a return to local, human, and transparent partners.

🚀 Launching Operation "Dindon!" 📢

But talk is cheap. It's time for action. Dindon and his association have announced the launch of the largest operation in European digital history: **Operation "Dindon!"**. The concept: the greatest data repatriation of the century. The objective: to organize, plan, and provide technical and financial support for the massive repatriation of data from thousands of European companies — moving it out of the Public Cloud and into Sovereign and local Clouds.

📜 The Historical Tribute (2011)

Why the name "Dindon!" with that angry exclamation point? It is a vibrant tribute to a courageous Director at STIME (the IT arm of the Mousquetaires Group) who, back in 2011, faced with the exorbitant costs of Windows Server and Office licenses, decided to name his major infrastructure migration project: "Operation Dindon." The circle is complete.

> 🌟 *Dindon's Final Word: The Move-to-Cloud is not a fate. The Move-to-Local is a liberation. Dindon has proven that you can be high-performing, profitable, and sovereign. The story ends here, but the movement is just beginning.*

👉 *What about you? Are you going to remain a "Cloud Fall Guy," or will you join Operation "Dindon!" to bring your data back to a partner you can trust?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode VII

The "Managed" Mirage — The Phantom Kubernetes Bill

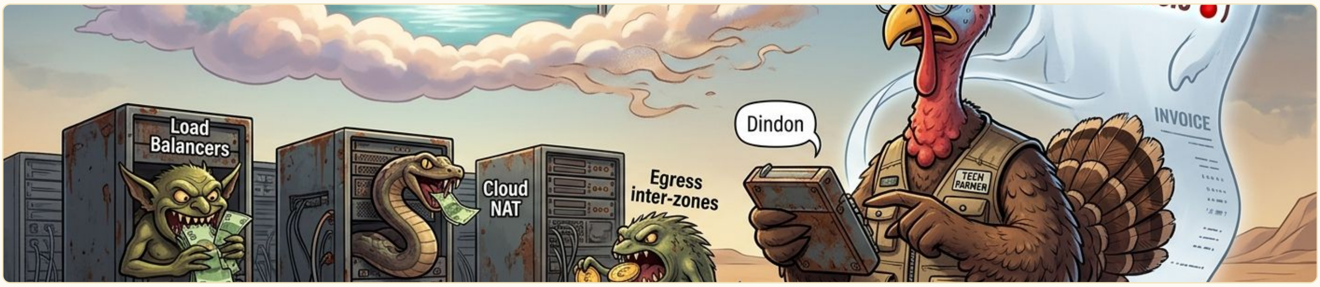

HELLO NETWORK! 👋

Remember Dindon? Our tech farmer who had fled the Public Cloud at €475/VM to find technical and financial sanity (sovereign Bare Metal, real ROI). Well... Dindon relapsed. 😔

He had kept one last addiction hidden in a corner of his console: his "Managed" Kubernetes cluster. He had convinced himself: "It's managed, it's simple, it's just the modern compute layer, it won't bankrupt me." Bullshit. 🛑 "Managed" doesn't mean "Free" — it means "Opaque."

📊 The Phantom Cluster Autopsy

Expense item	Imagined budget	Actual invoice
Control Plane (GKE Fee)	«Free?»	€72/month
4 Nodes (n2-standard-8)	~€600/month ✅	~€600/month ✅
Load Balancers (×3 forgotten)	«Negligible»	€54/month
Persistent Volumes (500 GB)	«Included?»	€85/month
Cloud NAT (Private cluster)	«Security...»	€45/month
Cloud Logging / Monitoring	«By default»	€120/month
Cross-zone Egress	«A few GBs»	€90/month
Ghost Nodes (Autoscaler)	«It's magic»	€160/month
MONTHLY TOTAL	**~€600** 💭	**~€1,226** 🩸

Bottom line: 104% in hidden extra costs.

🛑 The Moment of Truth

◆ His "managed" cluster: €14,712/year.

◆ These workloads would have easily fit on two Grade A refurbished dedicated servers. Cost: €4,000 per unit. Total: €8,000, a one-time hit (CAPEX). Negative ROI by the 7th month.

The worst part? "Managed" is just an Ops illusion. They manage the control plane, but YOU are still the one wrestling with complex drop-down menus for NAT, LBs, and volumes at 3:00 AM.

> *"Managed for them. Complicated for you. Expensive for everyone."*

⚖️ The Architect's Verdict

Render unto Caesar what is Caesar's: hyperscalers invented brilliant abstraction. But they also invented financial vassalage. True #FinOps means reading the bill line by line, not blindly trusting the abstraction that suffocates you.

👉 *What about you? Have you ever audited every single item of your "managed" cluster — including the ones you never ordered?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode VIII

The Highway to Hell — The AWS NAT Gateway Autopsy

Hello Network! 👋

Thought Dindon hit rock bottom with Kube? 🛑 Today, we tackle the ninja of financial asphyxiation: **The AWS NAT Gateway**.

For "High Availability," the AWS manual advises 1 NAT per zone (AZ). Dindon has 3 AZs, so he clicked 3 times. 🖱️ What he didn't know? At the hyperscaler's, doing NAT (a simple iptables rule, free on Linux since 1998) isn't a feature — it's a luxury tollbooth.

📊 Network Toll Autopsy

For 10 TB/month of outbound data (updates, APIs, webhooks...):

Expense (NAT Gateway ×3)	AWS Rate (Paris)	Actual Invoice
Existence Tax	$0.048 / hour	~€105/month
Processing (10 TB / month)	$0.048 / GB done	~€460/month
MONTHLY TOTAL	**«It's native»** 💭	**~€565** 🩸

🛑 The Moment of Truth

€565/month. €6,700/year. Just to change a TCP source IP! ☕💥

On-Premise, a €150 Mikrotik router (one-time CAPEX) or a pfSense VM (€0) handles this without breaking a sweat. In the Cloud, the network is a taxi meter spinning with every byte breathed.

The giants' "High Availability" is often designed to multiply your paid resources. True #FinOps? Replace those 500-buck managed NATs with a simple EC2 instance (NAT instance) or a dedicated local server for the price of a subway ticket. Or better yet: #CloudExit.

👉 *What about you? Have you ever dared to look at how much the simple MASQUERADE rule of your clusters costs you?* 😉

Amine RAITI · CC BY-NC-SA 4.0

Episode IX

Awakening the Sleepwalkers

Hello network! 👋

Today, we're dotting the "i"s and putting instances on real hardware. We're going to talk about a touchy subject, but we'll do it with respect: bankruptcy by the Cloud and the great sovereign slumber.

Dindon has a confession to make. If he has become an ayatollah of FinOps and Bare-Metal, it's not out of ideology. It's because he has seen ships sink. I have seen with my own eyes how infrastructure can go from being an "innovation engine" to a "financial ball and chain."

🌪️ The Hyperscaler: The Masterpiece That Ruins You

Let's be fair: technically, what AWS, GCP, or Azure do is masterful. Their engineers are geniuses. But their economic model is a cash vacuum designed for unicorns doped on fundraising. The tragedy is that SMBs and traditional companies are copying these architectures. The result? Public Cloud becomes the leading cause of financial stress for CFOs. We no longer go bankrupt because we lack clients; we go bankrupt because our variable costs (Egress, NAT, API calls) grow faster than our revenue.

🛌 Local Actors and Politicians: The Sleepwalkers' Club

And while the real economy gets siphoned by untraceable OPEX bills, what are our champions of "Digital Sovereignty" doing? They are sleeping. 😴 Our politicians make magnificent speeches with the phrase "Sovereign Cloud" in size 72 font. But on the ground? No massive tax incentives to relocate data, no real education on Total Cost of Ownership (TCO).

And our local hosters? You have the steel. You have the datacenters. You have the prices (sometimes 10× cheaper in CAPEX). But where is your voice?

> *O sleepers perched upon the cloud,*
> *Your pockets plundered by the crowd!*
> *You yield our iron to the foe,*
> *Now weep in blood as down you go.*

⚖️ The Moment of Truth

Public Cloud is not a scam; it's a luxury product. If you need instant global elasticity, pay the premium. But for 80% of workloads (classic databases, stable backends, web), staying on a hyperscaler is like using a private jet to go buy a baguette. True "Scale" is profitability.

👉 *Local hosters: wake up, we need you on the battlefield, not in VIP lounges.*

👉 *CTOs & CFOs: sit down together. Infrastructure is no longer a technical topic; it is a matter of financial survival.*

Amine RAITI · CC BY-NC-SA 4.0

Episode X·1

The Pride at the Top (Part 1/2)

Hello, network! 👋

Today, a change of program. Before judging local hosters or the State, Dindon has decided to audit the root of the evil: the company itself. Because the sinking of the Cloud is never a technical accident; it is always an organizational suicide.

Welcome to the autopsy of an internal disaster, a clinical pattern that our industry repeats in loops.

🎭 The Theater of "Cost Reduction"

Imagine a management committee organizing a major "workshop" to "rationalize the infrastructure." On the table, the numbers are lunar: hundreds of thousands of euros burned each month with hyperscalers, coupled with six-figure bills for "modern" SaaS licenses.

Yet, in the shadow, this same company possesses its own treasure: physical racks in local Datacenters, half-empty, perfectly equipped, and already amortized, costing ten to twenty times less to run.

💡 The Affront to Logic

Faced with this absurdity, the field Ops proposes the obvious: let's use the empty space of our Bare-Metal to repatriate our Kubernetes clusters and divide the production bill. The reaction from the top? Scorn. The CEO, blinded by the ego of his own "vision," brushes the proposal aside with a wave of his hand: "We'll see later, let's stay on the roadmap." The CTO, for his part, is silent — he is the archetype of the "Technical Director" who clings to derisory micro-tasks to justify his place, but who validates without batting an eye the waste of millions in OPEX.

📁 The Wall of Silence (and the Warning)

The field engineer does not give up. He produces a quantified study, multiple exit scenarios, quotes in support. The response? Management plays for time. For a whole week, his requests for a simple 15-minute slot are dodged. To force the door of his own management committee, the Ops literally has to slam his fist on the table and demand, by a frontal email in front of other employees, an emergency meeting.

The meeting is finally wrested, but the sanction falls: a recasting and an oral warning from management. Laying numbers on the table is no longer seen as financial optimization; it is perceived as insubordination. The accounting truth has become the sworn enemy of managerial pride.

Amine RAITI · CC BY-NC-SA 4.0

Episode X·2

The Silence of the Trenches (Part 2/2)

Hello, network! 👋 (Continuation of the internal audit — 2/2)

When the top of the pyramid (CEO, CTO, CFO) chooses blindness by ego, what does the rest of the company do in the face of the financial wall?

📢 The General Alert

Faced with denial, our field Ops decides to break the law of silence. He takes the debate to the internal public square: emails, meetings, face-to-face. He warns that the ship is heading straight into the iceberg. The objective is not to make a publicity stunt for his ego, but to save the profitability of the company and, by extension, the jobs of his colleagues.

🤫 The Cowardice of the "Cushy Jobs"

At first, colleagues encourage behind the scenes. But from the first management recasting email, it's a stampede. All departments withdraw from the debate. Support is whispered in secret, at the coffee machine. Faced with the risk of confronting management, the fear of losing their "cushy job" prevails over the instinct of collective survival. The whistleblower is left to isolate himself alone on the battlefield.

⚖️ The Veracious Witness

The engineer then makes a choice: to take the oath to remain the "veracious witness" of this intellectual bankruptcy, even if it must cost him his place. Before the ego of management and the complaisant silence of employees, he is inevitably pushed to the exit. The inconvenient witness is eliminated, Cloud waste continues, and colleagues keep their "cushy job"... waiting for the next social plan.

> *The boss wallows in error,*
> *The court applauds out of terror.*
> *You prefer to keep your places,*
> *While the gold melts in the ice.*

> ⚖️ *The Field Verdict: Do not complain about GAFAM bills if you do not have the courage to contradict a disconnected CEO or a straw CTO. The hyperscaler is not responsible for your ruin. Your pride at the top and your collective cowardice at the base manage that very well all by themselves.*

Amine RAITI · CC BY-NC-SA 4.0

🇬🇧 Episode XI

The Duty of Truth (Local Audit)

Hello network! 👋

The testimony of Operation Dindon now turns toward our local players. This is not about blaming the giants for their excellence, but about taking a sincere look at our own failings. Sovereignty cannot be a mere slogan; it must be a promise of competence.

🤝 Unity: A Shared Responsibility

The market follows efficiency. The call for a common front — ironically named "Operation Dindon" to trigger a wake-up call — was an invitation to protect our SMEs. Ignoring this extended hand is not just a management error; it is a failure in the duty of economic fraternity. To remain silent while our companies bleed financially is to become complicit.

💻 The Justice of Experience (DX): Respecting the Craftsman

The developer is the one who prescribes. Offering them dated portals and obscure documentation is an injustice. Hyperscalers respect their time with fluid APIs. Technical justice starts here: if we want a return to local iron, let's make the path passable instead of a calvary.

🔬 Sincere Engagement vs. The Mirage of the Label

The market wants managed services, not just raw metal. Leaving our companies with no alternative face-to-face with proprietary lock-in is a failure of counsel. Similarly, brandishing labels to mask a technical lag lacks sincerity. Patriotism must not be a shield for incompetence, but a driver of excellence.

⚖️ The Verdict

Our sovereignty sleeps in our datacenters, locked by a lack of audacity. No one will save the local sector in its place. The courage to unite and the sincerity to recognize its lags are the only keys.

> *Vendors of iron, find your art once more,*
> *Sovereignty is forged on every shore,*
> *When local knowledge finally takes the lead,*
> *The steel awakes and sows its own seed.*

Amine RAITI · CC BY-NC-SA 4.0

🇬🇧 Episode XII

The Mirror of the State (National Audit)

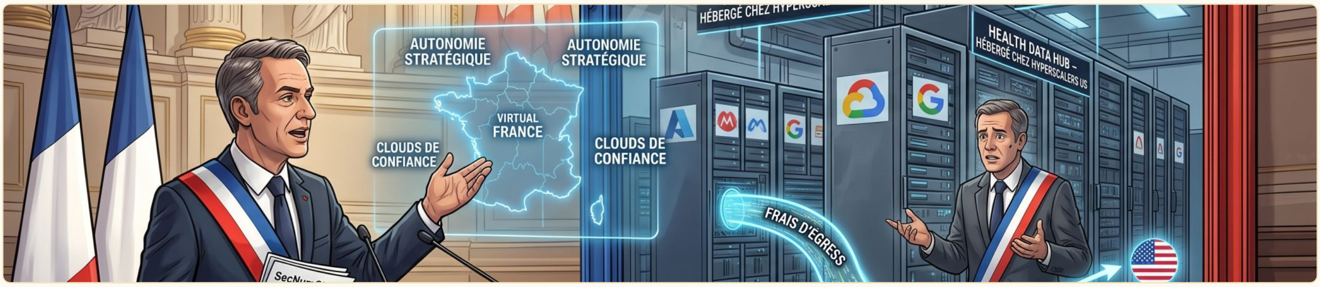

Hello network! 👋

The Operation Dindon audit rises today to the highest levels of the State. After analyzing the responsibilities of the local sector, the demand for justice compels us to examine public action. This is not about building a biased prosecution, but about placing the mirror of reality in front of stated ambitions.

🏛️ Rhetoric vs. Reality

Strategic autonomy is a noble ideal, fiercely defended in official speeches. However, the physics of Datacenters does not feed on eloquence. Declaring sovereignty while neglecting massive investment in hardware infrastructure creates a gaping dissonance. The truth of the Iron demands concrete actions, at the risk of seeing ourselves reduced to the status of digital tenants.

🏷️ The Paradox of the "Trusted Cloud"

Affixing a local seal (SecNumCloud) on a technology whose engineering and kill switch are located across the Atlantic is a legal construct, not technological independence. Sincere counsel forces the question: how can we demand our SMEs repatriate their data to sovereign infrastructures when the State itself entrusts its crown jewels (like the Health Data Hub) to hyperscalers? Authority is built on example, not on exemption.

💸 The Hemorrhage of Public Effort

The State legitimately waters our startups through generous subsidies (BPI, French Tech). Yet, in the absence of FinOps guardrails, these millions of euros too often end up absorbing exponential Cloud bills, de facto financing foreign R&D. Where are the tax incentives to encourage the return to CAPEX and investment in local Bare-Metal? The silence on this economic lever is deafening.

⚖️ The National Verdict

Sovereignty cannot be bought with rhetoric; it is forged in the steel of our own infrastructures. Legislating against data lock-in (Egress fees) and offering tax incentives for companies to return to Bare-Metal are the true levers of an assumed and lasting independence.

> *Beneath the soothing echo of sovereign speech,*
> *Public wealth drifts to a foreign shore.*
> *We must forge the steel within our reach,*
> *For true independence to live evermore.*

Amine RAITI · CC BY-NC-SA 4.0

🇪🇺 Episode XIII

The Continental Assessment (European Audit)

Hello network! 👋

This manifesto is drawing to a close. The French diagnosis (Episode XII) was merely the symptom of a broader failing architecture: the strategic flaw of the European Union.

🏛️ Legislative Asymmetry (The Brussels Paradox)

While Europe discourses on compliance and the EUCS label, a financial asymmetry takes hold. The European legislator managed to regulate telecom "Roaming" down to the cent, yet looks away when data is shackled by prohibitive egress fees. The Data Act remains an abstraction in the face of the reality of irrevocable 36-month "commits." Europe thus lets its SMEs be bound by general terms and conditions of sale that neutralize all technical freedom.

♻️ The Paradox of European Funds

Recovery plans and European funds (Horizon Europe) legitimately irrigate our technological fabric. Yet, the financial mechanics often end up draining this capital towards the San Francisco Bay. By refusing to impose local infrastructure quotas to obtain these subsidies, Europe indirectly finances its own technological trade deficit.

⚠️ Inertia as a Strategic Choice

Persisting in this legislative inertia, after being alerted to the reality of TCO and proprietary lock-in, changes the nature of the failure. This is no longer administrative sluggishness; it is accepting the dependence of our engineering by omission.

> *Brussels debates beneath its ceilings vast,*
> *While Egress drains our forces in silence.*
> *Sovereignty is born of Iron, forged to last,*
> *And engineering alone shall dictate the balance.*

🛑 The Withdrawal of the Witness (Return to the Forge)

This concludes the Operation Dindon audit. To the media, lobbyists, or podcasters: the media spectacle and sterile debates are politely declined. The withdrawal from this public circus is definitive. However, the posture of the "truthful witness" remains. The technical fight will continue where it began: in the engineering trenches, at the heart of the Bare-Metal, facing the Radix.

🔓 Open-Source Legacy

The entirety of the "Operation Dindon" manifesto (Episodes I to XIII) is released into the public domain. These texts may be reused, shared, or adapted by anyone, without authorization. Technical truth has no owner, and the urgency is now collective.

Amine RAITI · CC BY-NC-SA 4.0

🇬🇧 Episode XIV

The Seal of Sovereignty — Solemn Closure · Part 1/2

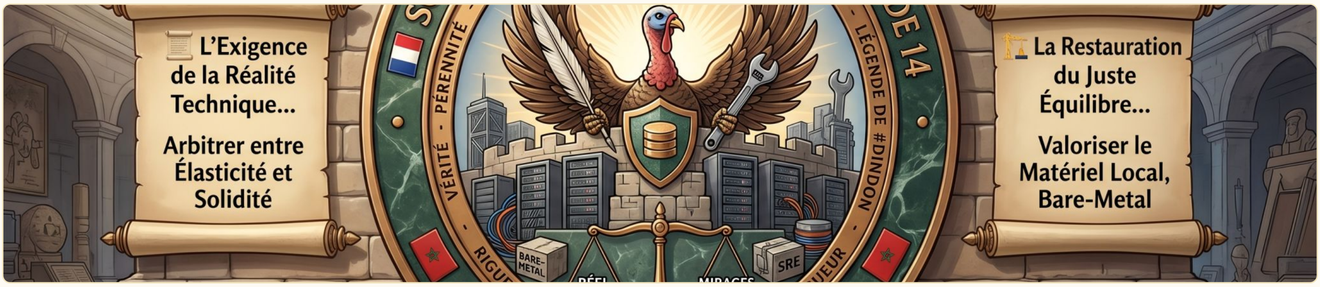

Hello network! 👋

The Operation Dindon diagnosis is coming to an end. After analyzing cost structures, strategic dependencies, and political levers, the time has come for a technical assessment and the closing of this chapter. This final episode is intended as a solemn synthesis, driven by the demand for truth and sustainability.

🏛️ The Demand for Technical Reality

Infrastructure is not just an accounting concept or a software abstraction; it is the physical foundation of our economies. The current trend towards total outsourcing has revealed its limits: loss of control over operational costs, increased complexity of transfers, and an erosion of field skills. The sustainability of a company rests on its ability to accurately arbitrate between service elasticity and asset solidity.

🏗️ The Restoration of the Right Balance

Digital sovereignty cannot be a mere slogan. It is built on precision engineering, valuing local hardware and the optimization of existing resources. The philosophy of returning to the "Iron" (Bare-Metal) and to technological recycling is not a step backward, but a step forward towards sounder and more responsible management. Strategic autonomy begins where dependence on infrastructures whose variables we no longer control ends.

Amine RAITI · CC BY-NC-SA 4.0

📜 The Mu'allaqa of the Final Seal

Ode to Rigor — Solemn Closure

> *Let tears for squandered budgets now be past,*
> *To face the long-forbidden roads at last.*
> *No cloud salvation without earthly might,*
> *For technical reason is our only light.*
> *Honor the structure, the assembly complete,*
> *Where each iron node makes our freedom concrete.*
> *The Balance is fixed on the weight of the real,*
> *Far from the noise and the mirrored ideal.*
> *Let the forge quiet down and let order endure,*
> *Justice is brought to the machine, swift and sure.*
> *Knowledge is a seal, and courage a proof,*
> *To build up a destiny under one's roof.*
> *The audit concludes, the verdict is cast,*
> *Restoring to Bare-Metal the pride of its past.*
> *Silence returns to this data array,*
> *History will judge the works ordered today.*

⚖️ The Audit's Verdict

Sovereignty is a choice. If I was harsh with companies, it is because their responsibility is decisive: neither the State nor Europe will act for them. As an employee myself, I know that the survival of our tools depends on our own audacity. It is now up to organizations to transform these findings into lasting decisions. Operation Dindon ends here, giving way to silent and rigorous action at the heart of the systems.

🦃

The Legend is complete.

The Operation continues.

Amine RAITI · CC BY-NC-SA 4.0

🦃

The Legend is complete.

The Operation continues.

◆ COLOPHON

Title: The Legend of Dindon
Author: Amine RAITI — Infrastructure Architect & SRE
Origin: XIV episodes published on LinkedIn, 2026
Collected and formatted: June 2026, as part of Operation Dindon
Licence: CC BY-NC-SA 4.0 — Attribution, Non-Commercial, Share Alike

This collection may be freely shared, reproduced and adapted for non-commercial purposes, provided the author is credited and any derivative work is distributed under the same licence.

AI Powered by Amine — AI is an amplifier of ideas and forms, not the source.

◆ OPERATION DINDON — CONTEXT

Operation Dindon is a personal, independent, public interest initiative launched on 8 May 2026 to document and challenge the predatory contractual practices of American hyperscalers (Amazon Web Services, Google Cloud, Microsoft Azure).

Three demands to the trio, deadline 8 September 2026:

- ◆ Cancellable commits
- ◆ Reasonable egress fees
- ◆ Real portability

The Legend of Dindon prefigures this fight. It is its narrative origin.

NEMO SUPRA LEGEM EST

وقيل الحمد لله رب العالمين

LETTER
OF A WITNESS

Open letter to His Holiness Pope Leo XIV
on the practices of the three pharaohs of the digital age

To

Sa Sainteté · His Holiness

Pope Leo XIV / Papa Leone XIV

Bishop of Rome · Servant of the Servants of God

From

AMINE RAITI

Infrastructure Architect & SRE
International Human · Universal Voice

Paris, 27 May 2026

The Duty of the Witness

I

Your Holiness,

I am not a Christian. I am a Muslim, an engineer, and a witness.

It is in this last capacity that I write to you — not to seek your blessing, nor to shelter behind your authority, but because the duty to bear witness for justice recognises no confessional border, no hierarchy of suffering, no calculation of opportunity.

> An-Nisa · Verse 135
>
> *"O you who have believed! Be persistently standing firm in justice, witnesses for Allah, even if it be against yourselves or your parents and relatives. Whether one is rich or poor — Allah is more worthy of both. So follow not your passions, lest you deviate. And if you distort your testimony or refuse to give it, then indeed Allah is ever, with what you do, Acquainted."*

This verse does not ask me whether the moment is favourable. It does not ask whether I have the means. It says: **bear witness**. Even against yourself. Even alone. Even when you tremble.

I bear witness.

✦

Your encyclical *Magnifica Humanitas* (15 May 2026) describes with prophetic precision what I have seen with my own eyes, counted with my own hands, and documented in my own invoices.

The main drivers of development are private, often transnational actors, endowed with resources and intervention capacities superior to those of many governments — a power essentially private, and therefore all the more difficult to identify, regulate, and direct toward the common good.

I submit their names:

Amazon Web Services · Microsoft Azure · Google Cloud Platform

These three entities are not tools. They are capture systems deliberately designed to transform dependency into debt, migration into punishment, and freedom into a priced privilege. They are the **pharaohs of our digital age**. Like the pharaohs of Egypt, they carry no whip — they carry contracts.

What I Have Seen

II

I have seen organisations forced to sign **36-month commitments**, transforming a technical decision into contractual servitude. When needs change, the exit is priced as a transgression. The contract does not serve the client. It holds them captive.

✦

I have seen data — *the data of these peoples* — taxed on exit at documented rates of **+9,931%** above market standards. Data does not belong to those who produced it. It belongs to those who hold it. And those who hold it charge for freedom.

✦

I have seen technical architectures designed not to serve the user, but to render them incapable of leaving. Proprietary services, incompatible formats, cascading dependencies — chains built in code, invisible to the naked eye, but as real as iron bars.

✦

This is not negligence. It is a **deliberate architecture of capture**. And those who suffer are not abstract corporations — they are engineers who lose their jobs when the invoice doubles, hospitals that cannot migrate their health data, governments that have surrendered their digital sovereignty.

"Technology takes on the face of those who design it, finance it, regulate it, and use it."

The face of these three actors is not that of service. It is that of **extraction**.

Name Them

III

You have condemned the technocratic paradigm that "*reduces creation to an object of exploitation and persons to cogs in a system that must be made ever more efficient.*"

I submit that this paradigm now has addresses, registration numbers, stock prices and PR teams. The Church has known, in other times, how to name slavery by its name — even when slaveholders were friendly powers and commercial partners. That courage cost it enemies. It also earned it its soul.

I ask for that same courage today.

Not to join my struggle — I seek neither institutional ally nor confessional endorsement. My independence is non-negotiable.

I ask you to name these practices for what they are: a contemporary form of economic enslavement. Name them. Not for me. For the millions who suffer without the words to say so.

❦

I do not know whether this letter will reach you.

But I write it, because verse 4:135 gives me no choice. The witness who evades their testimony through calculation or fear becomes complicit in the lie they allow to prosper.

I have counted the bytes. I have read the contracts. I have published the figures. I have composed songs in twenty-eight languages. I have waged this asymmetric war at fifty euros a month.

I am not a saint. I am an engineer who refused to look away.

With respect, lucidity, and the peace we share in the pursuit of justice,

AMINE RAITI

Infrastructure Architect & SRE

Initiator — Operation Dindon

CC BY-NC-SA 4.0 · 3.14

OPÉRATION DINDON · INTERNATIONAL EDITION

PART XI

SYNTHESIS

2 pieces

HUMAN

OPÉRATION DINDON · INTELLECTUAL SYNTHESIS · JUNE 2026

ORIGINAL THEORIES AND CONCEPTS

What the corpus discovered and named for the first time

◆ CONTEXT OF THIS SYNTHESIS

The Opération Dindon corpus now comprises more than twenty structural studies. Beyond factual analyses and public policy proposals, this corpus has produced original concepts and theories — new words, mechanisms not yet formalised, theses not yet stated in these terms. This synthesis inventories and describes them, so that their significance is not lost in the volume of the corpus.

NEOLOGISMS	THESES	FORMULAS
3	9	10

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Infrastructure instructor

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

THREE WORDS THAT DID NOT EXIST

◆ 1 — INFRATIONAL / INFRATIONNELLE

Source study: The Infrational Crisis (2026)

The progressive devaluation of technical words through overuse — exactly as monetary inflation dilutes the value of currency. When "infrastructure" can mean a VM created by Terraform, when "DevOps" can mean a position halfway between two distinct professions, words stop protecting the realities they once named. And when words give way, skills follow. The neologism fills a documented lexical gap: the phenomenon existed, it had no name yet.

◆ 2 — THE INFRATIONAL LOOP

Source study: The Infrational Loop (2026)

The self-sustaining six-step cycle: DevOps dilutes bare-metal competences → profile shortage → hiring salaries explode (8 to 12%/year) → internal salary compression (2 to 4%/year) → seniors leave or take voluntary gap → shortage deepens → cycle restarts. Nobody at the wheel. Each actor acts rationally. The system is irrational. The loop closes without any actor having wanted this result — and without any isolated actor being able to interrupt it alone.

◆ 3 — THE FINOPS OF THE LAST GRAM

Source study: The FinOps of the Last Gram (2026)

The extension of FinOps to the physical residual value of end-of-life hardware — recovery of gold (€17 to €42/server), copper, palladium. Traditional FinOps stops at the cloud invoice. The FinOps of the Last Gram begins where the server ends — at the atomic level of the metal. A rack of 40 servers contains €880 to €2,300 in raw metallurgical value that the WEEE chain absorbs in its margins.

CORPUS

Amine RAITI · CC BY-NC-SA 4.0

MECHANISMS NOT YET FORMALISED

◆ 4 — THE NEUROCHEMICAL RECHARGE SEQUENCE

Sources: Play in SRE · The Body First · The Great Recharge

A complete recovery architecture in five frequencies — weekly Nerf (€0, fiero + adrenaline) → monthly dinner (oxytocin) → quarterly hammam (€35, cortisol ↓) → biannual activity (fiero + collective emotion) → annual Great Recharge (€350-500, ART/SRT, autonomy). Each format produces a distinct and documented neurochemical mechanism. No existing model combined these five levels in a coherent architecture at €600-900/year/person.

◆ 5 — THE SERVER ROOM AS MEDINA

Source: *The Mentor and the Metal* (2026)

The structural parallel between the Maâllem/Metaâllem process of Moroccan ancestral arts (Zellij of Fès, Gibs, Marrakech leather) and tacit knowledge transmission in bare-metal infrastructure. Both share the same survival mechanism: breaking the transmission chain kills the art or the profession. The Zellij did not survive because of books about Zellij — it survived because Maâllems accepted taking Metaâllems beside them.

◆ 6 — TACIT KNOWLEDGE AS THE INVISIBLE LAYER

Source: *The Mentor and the Metal* (2026)

The distinction between explicit knowledge (transmissible through training) and tacit knowledge in infrastructure — incident intuition, judgement under pressure, team culture. The senior engineer who enters the server room and knows, before looking at metrics, that something is wrong — this knowledge cannot be documented. It is forged alongside a Maâllem. Training without a mentor produces a technically trained but operationally incomplete engineer.

◆ 7 — THE OBSERVATION DAY AS MINIMAL FORMAT

Source: *The Mentor and the Metal* (2026)

The most accessible mentorship format: one day, not an internship, not an assessment. The student watches. They visit the racks. They eat with the team. They observe an incident if luck allows. This physical moment — this sensory encounter with infrastructure — produces an irreversible vocational trigger that six months of theoretical courses cannot produce. HR makes the connection with the school. The SRE does their work normally. That is all.

◆ 8 — THE DISSOLUTION OF THE BODY THROUGH VOCABULARY

Sources: The Infrational Crisis · The State and the Invisible Body · The Infrational Loop

The thesis that the disappearance of infrastructure competences was not caused by a direct decision but by a linguistic strategy — the progressive replacement of precise titles (systems administrator, SRE, network engineer) by vague titles (DevOps engineer, cloud engineer, IT coordinator). You cannot lose what you cannot name. You cannot name what you have lost the definition of. Word inflation preceded and conditioned the dissolution of skills.

CORPUS

Amine RAITI · CC BY-NC-SA 4.0

MECHANISMS NOT YET FORMALISED — CONTINUED

◆ 9 — DOCUMENTATION AS GENDERED NORMALISATION SIGNAL

Source: The Invisible Amputation (2026)

Analysis of name examples in AWS IAM documentation reveals a documentable structural bias: Jane, Alice, Adele appear in basic user creation and restricted permission roles — John, Bob, David dominate administrator roles, AssumeRole, CloudTrail. The canonical IAM security problem features John Doe creating the risk, Jane Doe inheriting it passively. This is not a conspiracy. It is a normalisation that enters the beginner engineer's cognitive unconscious. It may be more effective than a conspiracy.

◆ 10 — THE SRE AS STRUCTURAL ENEMY OF HYPERSCALERS

Sources: Anatomy of the Loss · The Invisible Amputation · The State and the Invisible Body

The SRE engineer who masters the physical layer is the structural — not personal — enemy of hyperscalers, precisely because they can always fall back to bare metal. Their competence is an existential threat to the lock-in economic model. A client who can leave is a client who negotiates. A client who cannot leave is captive. Reducing this body by 50% through women's exclusion simultaneously amputates it in numbers, in different perspective and in institutional legitimacy.

◆ 11 — THE NO-PROGRAMME PROGRAMME AS RESTORATION MECHANISM

Source: The Great Recharge (2026)

The thesis that the absence of a programme in a retreat is the central neurological mechanism of restoration — not an organisational oversight. An imposed programme maintains directed vigilance: I must be in the right place at the right time, do the right activity. This vigilance keeps the sympathetic nervous system on alert. The total absence of programme frees the brain from all demands. Deci and Ryan (Self-Determination Theory) document that perceived autonomy reduces cortisol independently of the activity chosen. Freedom is the programme.

◆ 12 — SYSADMIN ALCHEMY AS PHYSICAL PEDAGOGY

Source: The FinOps of the Last Gram (2026)

Recovering precious metals from end-of-life hardware is not only a FinOps exercise — it is a pedagogical format. The engineer who has melted their gold viscerally understands why gold fingers exist, why PCIe connectors are gold and not copper, why a server contains real physical value. The chain TSMC → chip → server → recovered gold makes the Digital Iron thesis visible and tangible. It is the Foundation of Iron through chemistry.

CORPUS

Amine RAITI · CC BY-NC-SA 4.0

THE WORDS THAT STRUCTURE AND THE PHRASES THAT CONDENSE

◆ 13 — SYSOPS / NETOPS / OPINFRA — THE THREE-CATEGORY TAXONOMY

Source: The Ticket and the Talent (2026)
The first structured nomenclature for digital infrastructure roles in three functional categories and six hierarchical levels. SysOps: from metal to service (servers, virtualisation, storage). NetOps: connectivity (switching, routing, firewalls). OpInfra: operational transversality (SRE, architects). Six levels: Support → Technician → Administrator → Engineer → Lead/Principal → Architect. The founding rule: never go beyond this naming. DevOps is a philosophy — not a title in this nomenclature.

◆ 14 — THE NON-DILUTION RULE

Source: The Ticket and the Talent (2026)
If a title cannot be expressed in the words of the SysOps/NetOps/OpInfra nomenclature, either the profile is poorly defined, or the title was invented to mask a classification problem. The existing words suffice to describe any real infrastructure profile. Any undefined composite title — "Senior Cloud DevOps Infrastructure Expert", "SRE Platform Tech Lead" — is a signal of dilution, not expertise.

◆ FORMULAS — PHRASES THAT CONDENSE A THESIS ◆

*"Metal comes before the code.
Metal comes before the cloud.
Metal comes before the promise."*
— The Foundation of Iron

*"You cannot lose
what you cannot name."*
— The Infrational Crisis

*"The blacksmith who does not know
their metal forges wind."*
— The Foundation of Iron

*"The server room is a medina.
The cables are the alleyways.
The racks are the workshops."*
— The Mentor and the Metal

*"The Zellij did not survive
because there were books
about Zellij."*
— The Mentor and the Metal

*"Saving €10,000
to lose €100,000."*
— The Infrational Loop

*"What is missing
is the decision."*
— The Great Recharge

*"The Nerf gun was not a toy.
It was a mental maintenance
instrument."*
— Play in SRE Teams

*"DevOps is a philosophy.
Not a profession."*
— The Ticket and the Talent

*"The neurochemistry is the same.
The budget is not."*
— The Body First

These concepts are freely usable under CC BY-NC-SA 4.0. Attribution required: Amine RAITI — Opération Dindon.

NEMO SUPRA LEGEM EST

CORPUS

ٱلْحَمْدُ لِلَّهِ وَسَلَامٌ عَلَىٰ عِبَادِهِ ٱلَّذِينَ ٱصْطَفَىٰ

MAGISTERIAL STUDY · COMPLETE CORPUS SYNTHESIS · SINCE 8 MAY 2026

OPÉRATION DINDON

Complete Map of a Corpus

Produced in a Few Days of Intensity

◆ **WHAT THIS DOCUMENT IS**

This document is the map of the Opération Dindon corpus — the complete set of studies, analyses, manifestos, tutorials and tools produced since 8 May 2026, in a few days of intensive work. It does not replace the studies — it shows how they articulate around a single founding thesis and reinforce each other. Eight thematic arcs. One central thesis. One ultimatum. One date.

STUDIES	ARCS	LANGUAGES	ULTIMATUM
25+	8	3	8 SEPT 2026

Amine RAITI — Infrastructure Architect & SRE

Former engineering school professor · Teaching since 2006

Public document · CC BY-NC-SA 4.0 · AI Powered by Amine · Opération Dindon

NO SOVEREIGNTY WITHOUT MATTER

◆ THE THESIS IN ONE SENTENCE

The human is sovereign when they can act on the matter they depend on.

Purely extraterritorial cloud removes this capacity.

Bare-metal restores it.

Everything else is consequence.

This thesis — explicitly stated in "No Sovereignty Without Matter" — is the keystone from which all corpus studies are extensions. It was not stated first. It emerged as a retrospective self-evidence after the corpus had documented its consequences from twenty different angles. That is the characteristic of a true thesis: it was there from the beginning, implicit in every study, before being named.

◆ THE THREE LEVELS OF THE THESIS

Philosophical: Bodin, Schmitt, Hobbes converge on the same principle — no sovereignty without control of real territory. The landowner-farmer is sovereign. The tenant is precarious. Applied to the digital world: the organisation that masters its bare-metal is sovereign. The one that rents its infrastructure from an extraterritorial hyperscaler is precarious.

Technical: bare-metal is digital territory — cables, servers, racks, datacentres. Cloud extracts this territory from the control of its dependant across four documented dimensions: unknown physical location, extraterritorial jurisdiction (CLOUD Act), unilateral termination, pricing power.

Human: technical sovereignty rests on a professional body capable of building and operating independent infrastructure. This body — bare-metal SREs, systems administrators, network engineers — is the human condition of sovereignty. Without it, any sovereignty policy remains nominal.

RATIO

◆ THE GENERAL PRINCIPLE THE CORPUS DISCOVERED

What resists abstraction is precious. Bare-metal resists cloud abstraction. Human lived experience resists AI abstraction. Tacit knowledge resists the abstraction of theoretical training. In every domain where abstraction takes hold, the value of what it cannot capture rises. This principle — discovered across eight different thematic arcs — is the corpus's central intellectual contribution.

Amine RAITI · CC BY-NC-SA 4.0

HOW THE CORPUS ESTABLISHED THE DIAGNOSIS

The corpus began with the diagnosis. Before proposing solutions, it was necessary to name with precision what is happening — how hyperscalers capture organisations, how vocabulary dissolves competences, how digital geography determines sovereignty.

◆ ANATOMY OF THE LOSS — THE CAPTURE MECHANISM

The corpus's first study documents the six lock-in layers: technical (proprietary APIs, non-portable formats), cognitive (certifications that chain teams), contractual (asymmetric termination clauses), economic (egress fees), cultural (cloud as implicit norm), and political (absence of binding regulation). The central thesis: lock-in is not an accident — it is the product architecture.

◆ THE INFRATIONAL CRISIS & THE INFRATIONAL LOOP — THE NEOLOGISMS

Two original neologisms. *Infrational*: the devaluation of technical words through overuse — "DevOps", "cloud engineer", "XxxOps" — which dissolve competences by making them indistinct. *The Infrational Loop*: the self-sustaining six-step cycle — dilution → shortage → salary explosion → compression → seniors leave → shortage deepens → restart. Nobody at the wheel. The collective result is irrational. Each actor acts rationally.

◆ THE DIGITAL GROUND & THE DIGITAL IRON — THE GEOPOLITICS OF THE SUBSTRATE

The Digital Ground documents datacentre geography — the AMS-FRA-DUB concentration, Interxion's acquisition by Digital Realty, the CLOUD Act vs GDPR mapping. *The Digital Iron* goes further: TSMC, ASML, Chinese rare earths — the chip is the ingot of the 21st century. Digital sovereignty begins in the mines, not the datacentres.

◆ THE ECONOMY OF SILENCE & THE INVISIBLE DEBT

The Economy of Silence: uptime as the only metric — what the silence of the server room hides about the SRE engineer's real work. *The Invisible Debt*: the technical debt accumulated by organisations that migrate to cloud without mastering what they abandon. Both documents document what dashboards do not measure — yet which determines real resilience.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

THE PROFESSIONAL BODY AS LIVING TERRITORY

Digital sovereignty cannot exist without engineers capable of making it operational. Arc II documents this professional body — its pathologies, its exclusions, its pay, its departures, and the institutional conditions for its survival.

◆ THE SILENCE OF THE ROOMS & THE INVISIBLE AMPUTATION — THE MISSING HALF

The Silence of the Rooms documents women's absence in SRE and infrastructure: 10 to 15% women in French SRE in 2025, against 57% in medicine. *The Invisible Amputation* goes further: analysing name examples in AWS IAM documentation — Jane, Alice in restricted roles; John, Bob in administrator roles — as a gendered normalisation signal. SRE without women is a body amputated by half — less creative, less resilient, less institutionally legitimate.

◆ MATERNITY AS CONTINUITY & THE BLUE

Maternity as Continuity inverts the perspective: maternity is not a career interruption — it is a skills continuity (priority management under pressure, anticipation, resilience). *The Blue* documents what the diploma does not measure — craft knowledge, the culture of physical work, the tacit competence of despised vocational pathways. Both studies defend the same thesis: real value is not measured with standard instruments.

◆ THE DEPARTURE OF THE LAST ONE WHO KNOWS & THE STATE AND THE INVISIBLE BODY

The Departure of the Last One Who Knows documents the precise moment when an organisation irretrievably loses its tacit knowledge — and what it actually costs. *The State and the Invisible Body* poses the institutional question: why does the French state not recognise, train, and protect the professional body that is the condition of its own digital sovereignty? Six proposed levers: protected RNCP title, France Travail convention, public procurement preference, annual report, pay revaluation, mandatory Foundation of Iron in BTS.

◆ THE CPF MISUSED & THE TICKET AND THE TALENT

The CPF Misused documents how training rights are massively oriented toward hyperscaler cloud certifications at the expense of sovereign infrastructure training. *The Ticket and the Talent* proposes the correction: a structured nomenclature of infrastructure roles in three categories (SysOps/NetOps/OpInfra) and six levels, with the Non-Dilution Rule — DevOps is a philosophy, not a title.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE ECONOMICS OF BARE-METAL — FROM SERVER TO GRAM OF GOLD

Arc III documents the real economic value of bare-metal infrastructure — at entry (refurbished hardware), during life (cloud cost reduction), and at end of life (precious metal recovery). The corpus's FinOps does not stop at the cloud invoice.

◆ GRADE A REFURBISHED — ENTRY FROM BELOW

Grade A refurbished hardware is not a compromise — it is a strategy. Dell R640 at €800-1,200 versus €4,000-6,000 new, reconstructed manufacturer warranty, identical performance for the majority of SRE workloads. *Grade A Refurbished* documents the complete chain: sourcing, certification, integration, monitoring specific to reconditioned equipment.

◆ THE FINOPS OF THE LAST GRAM — EXIT FROM THE TOP

An end-of-life Dell R640 contains 0.2 to 0.5g of gold (€17 to €42), copper, palladium. A rack of 40 servers: €880 to €2,300 in precious metals that the WEEE chain absorbs in its margins. The FinOps of the Last Gram documents the complete DIY process — HCl/H_2O_2 streaming, $FeSO_4$ precipitation, NaOH neutralisation, torch melting. Lab budget: €245 amortised over 5 batches. Net batch 1: €780 to €2,200. Physical pedagogy as a bonus: the engineer who has melted their gold viscerally understands what is in a server.

◆ THE SRE CLOUD COST REDUCTION GUIDE — IN-LIFE OPTIMISATION

12 pages dedicated to cloud cost reduction for AWS/GCP/Azure SRE and Ops teams — rightsizing, reserved instances, spot instances, network optimisation, egress fee reduction, non-production shutdown automation. The corpus's most immediately operational document: applicable immediately, without infrastructure change, by the team in place.

◆ VMWARE → PROXMOX MIGRATION — OPERATIONAL RECONQUEST

Two versions (10p and 16p) of the VMware to Proxmox VE migration tutorial — Stretched Cluster D1/D2/D3, live migration, zero downtime. The VMware migration is the most accessible concrete sovereignty decision for French organisations following Broadcom's licensing policy changes. The tutorial documents every step, every command, every verification point.

RAITO

Amine RAITI · CC BY-NC-SA 4.0

MAINTAINING HUMAN BEINGS IN OPERATIONAL CONDITION

A rare and precious professional body that burns out is a body that disappears. Arc IV documents the complete recharge architecture — from the weekly micro-dose to the annual grand recharge — with its scientific basis and economic calculation.

◆ THE PAGER AND THE BODY — THE CORTISOL DEBT

On-call produces a documented accumulation of nocturnal cortisol. An SRE on regular on-call accumulates a physiological debt that holidays do not entirely repay. *The Pager and the Body* documents the mechanism, the biological markers, and the alert thresholds. It is the clinical basis on which the following studies build their solutions.

◆ THE YEMENI DINNER & PLAY IN SRE TEAMS — SHORT FORMATS

The Yemeni Dinner: the thoughtfully chosen restaurant as a signal of the manager's curiosity and life outside the office — €78 per person, bonding oxytocin, shared memory, real dehierarchisation. *Play in SRE Teams*: the Nerf gun was not a toy — it was a mental maintenance instrument. Fiero, adrenaline, dehierarchisation, micro-disconnection. These two formats document the monthly and weekly frequency of the Neurochemical Recharge Sequence.

◆ THE BODY FIRST — THE BARBÈS → BELLEVILLE SEQUENCE

Barbès hammam rented entirely (€150-200 for 10 people, 1h) → free walk to Belleville → Asian dinner. €35 per person. Four documented neurochemical phases: cortisol ↓ (hammam, Rapaport 2012), endorphins (walking, Tomasello/Wolf 2023), bonding oxytocin (dinner), respect signal (autonomous return). The freedom of route is the mechanism — not the destination.

◆ THE GREAT RECHARGE — THE DEEP ANNUAL

Mountain weekend at hot springs. No programme. No forms. No consultant. €350 to €500 per person. Kaplan (ART, 1989), Ulrich (SRT, 1983), Deci and Ryan (Self-Determination Theory), de Bloom et al. (22-study meta-analysis). Complete architecture: Nerf (€0/week) → dinner (€78/month) → hammam (€35/quarter) → activity (€75/semester) → mountains (€400/year) = €600-900/year/person.

HUMAN

Amine RAITI · CC BY-NC-SA 4.0

THE FOUNDATION OF IRON — FROM ELECTRON TO SYSTEM

Technical sovereignty is not decreed. It is transmitted. Arc V documents the mechanisms of tacit knowledge transmission and the concrete proposal for rebuilding the professional body — from the philosophy of mentorship to the complete training programme.

◆ THE MENTOR AND THE METAL — THE PHILOSOPHY OF TRANSMISSION

The structural parallel between the Maâllem/Metaâllem of Moroccan ancestral arts — Zellij of Fès, Gibs, Marrakech leather — and the transmission of tacit knowledge in bare-metal infrastructure. Both share the same survival mechanism: breaking the transmission chain kills the art or the profession. The Zellij did not survive because of books about Zellij. The Observation Day as minimal format: one day without assessment, without badge — the student watches, eats with the team, lives an incident if luck allows.

◆ THE FOUNDATION OF IRON — 26 WEEKS · FR AND EN · 52 MATERIALS

The most complete infrastructure training programme produced by the corpus. 26 weeks, from electricity to Kubernetes, designed for all audiences. Week 1: electricity and power. Week 6: PC and server hardware. Week 8: Linux fundamentals. Week 17: OSI model. Week 22: Active Directory. Week 26: network security and final defence. Each support exists in FR and EN — 52 A4 files, imperial format, print-ready. It is the operational answer to "how do we rebuild the professional body?" — starting from zero, with anyone who wants to learn.

◆ MASTERY OF THE IRON & THE TEACHER'S MANIFESTO

Mastery of the Iron completes the Foundation by documenting the advanced level — the final lock and the path to reconquest. *The Teacher's Manifesto* sets the posture: Amine does not speak as an institutional representative — he speaks as a trainer to his former students, with twenty years of classrooms behind him. "I have something more solid than a mandate: twenty years of classrooms."

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHAT CONTRACTS SAY — AND WHAT THEY HIDE

Digital sovereignty has a legal dimension that the corpus has documented with unusual precision for a technical corpus. The contractual analyses produced are real working documents — usable by CIOs, legal teams and CFOs in their negotiations with hyperscalers.

◆ AWS / AZURE / GCP TERMS ANALYSES — FR, EN, DE

Comparative analyses of the three hyperscalers' terms of service — in French, English and German — against French law, European law (GDPR) and post-Brexit British law. Points documented: unilateral termination clauses, liability limitations, extraterritorial data transfers, absence of binding SLA, unilateral modification of conditions. These analyses are the first in the corpus to show that cloud terms are not balanced contracts — they are instruments of legal captivity.

◆ THE CONTRACTUAL WAR MACHINE — DGCCRF, CMA, BUNDESKARTELLAMT VERSIONS

Adaptation of the contractual analysis to the three competition regulatory authorities — DGCCRF (France), CMA (United Kingdom), Bundeskartellamt (Germany). Each version identifies clauses potentially constituting anticompetitive or abusive practices before the relevant authority. These documents are designed to be submitted directly to authorities.

◆ THE CLOUD EXIT GUIDES — SHEETS A TO E

Five operational sheets dedicated to the five actors of a cloud exit decision: Sheet A (Legal Department), Sheet B (CIO — technical sequencing), Sheet C (CFO/General Management — ROI and hidden costs), Sheet D (complete operational sequence), Sheet E (HR and change management). These guides allow each actor to understand their role in a partial or full hyperscaler cloud exit.

◆ THE CONTRACTUAL NEGOTIATION GUIDE — 7 PAGES

Clauses to negotiate imperatively: egress fee caps, right to datacentre audit, minimum termination notice, guaranteed contractual portability, symmetric liability limitation, tariff freeze clause. This guide gives CIOs the concrete levers to negotiate from a position of strength — or confirms they have none, which is itself strategic information.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHAT ABSTRACTION CANNOT CAPTURE

The corpus opened an unexpected arc — the question of AI, generated content and human singularity. In dialogue with Sirine Houara, it discovered that the principle "what resists abstraction is precious" applies beyond infrastructure.

◆ FROM AMPLIFIER TO SUBSTITUTE — THE CRITICAL BOUNDARY

AI as amplifier: it takes ideas forged in experience and structures them, enriches them, gives them form. AI as substitute: it produces instead of the human, who only validates. The result may be stylistically correct — it is ontologically hollow. *From Amplifier to Substitute* documents the boundary and proposes criteria for staying on the right side.

◆ THE REPLACEMENT THAT REVEALS — IN DIALOGUE WITH SIRINE

Sirine stated the thesis: AI makes humans more precious by making them scarce. Lived experience is irreproducible — a life trajectory cannot be trained on. Contradiction is the sign of living thought — an LLM optimises for coherence, that is its weakness. Embodiment means the dyslexic's spelling errors are more authentic than the model's perfect text. The corpus established the link: same mechanism as the Infrational Loop — abstraction reveals the scarcity of what it cannot capture.

◆ AI AT THE SERVICE OF HUMANS — 11 PAGES

The complete training document on AI use in an infrastructure professional context. How to use LLMs without dissolving into them, how to integrate them into SRE workflows, how to maintain technical mastery in the face of growing automation. The training that completes the Foundation of Iron for the generation entering the profession with AI as a native tool.

◆ THE €21.99 AI NOTE AS LIVING DEMONSTRATION

Every corpus post ends with the same note: "🤖 Commissioned AI disclosure: produced under the instructions of Amine RAITI... Monthly salary: €21.99." This is not only humour. It is a transparency declaration — the tool is named, the author is identified, the amplifier/substitute boundary is visible. The theses of the corpus are irreproducible because they come from twenty years of field work. AI gives them their form. It does not give them their substance.

RAITI

Amine RAITI · CC BY-NC-SA 4.0

8 MAY 2026 · FOUR DEMANDS · 8 SEPTEMBER 2026

◆ THE ULTIMATUM — FULL TEXT OF THE FOUR DEMANDS

Issued on 8 May 2026 by Amine RAITI, addressed to AWS, Microsoft Azure and Google Cloud Platform:

1. Resiliable commits: multi-year commitments must be terminable without disproportionate penalty in cases of unilateral modification of conditions by the provider.

2. Reasonable egress fees: outbound data transfer charges must be capped at a level that does not constitute an economic barrier to mobility.

3. Real portability: data must be exportable in standard open formats, within reasonable timeframes, without technical degradation.

4. The CLOUD Act stays in the USA: decisions of the US Department of Justice do not apply to data of European citizens and organisations hosted in European datacentres.

◆ WHAT WAS PRODUCED BETWEEN 8 MAY AND 8 SEPTEMBER 2026

In a few days of intensive work since 8 May 2026: more than 25 structural studies FR+EN, 52 Foundation of Iron teaching supports, legal analyses of AWS/Azure/GCP terms in three languages, operational cloud exit guides, MariaDB and Proxmox technical tutorials, the trilingual Legend of Dindon, Latin lexicons, and this magisterial synthesis. The corpus is not the product of two years of progressive writing — it is the product of a few days of total intensity. This is not a weakness. It is a demonstration.

◆ WHAT REMAINS — 8 SEPTEMBER 2026

The ultimatum runs until 8 September 2026. If by then no concrete response — contractual, tariff-based, legislative — is provided to the four demands, production continues. This corpus is not a threat — it is documentation. What is documented is not lost. What is named can be defended. What is defended can be reconquered.

RATIO

Amine RAITI · CC BY-NC-SA 4.0

ORIGINAL CONCEPTS — INVENTORY

◆ NEOLOGISMS

Infrational / Infrationnelle — devaluation of technical words through overuse

The Infrational Loop — self-sustaining 6-step cycle, nobody at the wheel

The FinOps of the Last Gram — FinOps extended to residual metal value

◆ NOMENCLATURE

SysOps / NetOps / OpInfra — 3 categories, 6 levels

The Non-Dilution Rule — DevOps is a philosophy, not a title

The Neurochemical Recharge Sequence — 5 frequencies, €600-900/year

◆ ORIGINAL STRUCTURAL THESES

The Server Room as Medina — parallel between Moroccan Maâllem/Metaâllem and SRE tacit knowledge

Tacit Knowledge as the Invisible Layer — what training without a mentor cannot transmit

The Observation Day as Minimal Format — one day without assessment as a vocational trigger

The Dissolution of the Body Through Vocabulary — lost words precede lost competences

Documentation as Gendered Normalisation Signal — AWS IAM John/Jane as implicit encoding

The SRE as Structural Enemy of Hyperscalers — bare-metal competence is an existential threat to lock-in

The No-Programme Programme as Restoration Mechanism — absence of structure releases directed vigilance

SysAdmin Alchemy as Physical Pedagogy — melting gold to understand what is in a server

◆ THE CORPUS FORMULAS

"Metal comes before the code. Metal comes before the cloud. Metal comes before the promise." · "You cannot lose what you cannot name." · "The blacksmith who does not know their metal forges wind." · "The server room is a medina. The cables are the alleyways. The racks are the workshops." · "The Zellij did not survive because of books about Zellij." · "Saving €10,000 to lose €100,000." · "What is missing is the decision." · "DevOps is a philosophy. Not a profession." · "The neurochemistry is the same. The budget is not." · "What resists abstraction is precious."

RATIO

Amine RAITI · CC BY-NC-SA 4.0

WHAT ALL OF THIS SAYS

In one page

◆ THE THESIS

No digital sovereignty without mastery of bare-metal.

No mastery of bare-metal without the engineers who understand it.

No engineers without the training that forges them.

No training without the trainers who have lived what they teach.

Everything else is consequence.

◆ THE EIGHT ARCS IN ONE SENTENCE EACH

I — The Diagnosis: lock-in is not an accident — it is the product architecture.

II — The Invisible Body: the SRE professional body is the human condition of sovereignty — it is scarce, underpaid, undertrained, and half absent.

III — The Value: from the refurbished server to the recovered gram of gold — bare-metal has real physical value that nobody calculates.

IV — The Care: a body that burns out is a body that disappears — €600 per year per person is enough to maintain the neurochemical cohesion of an SRE team.

V — The Transmission: knowledge is not documented — it is transmitted body to body, in real conditions, alongside a Maâllem.

VI — The Legal Arsenal: cloud contracts are instruments of captivity — they can be negotiated, contested and used as evidence before regulatory authorities.

VII — The Human in the Machine: what resists AI abstraction is precious — the irreproducible lived experience, contradiction, embodiment.

VIII — The Ultimatum: four demands, one date, a corpus as testimony.

This corpus is not a threat. It is documentation.
What is documented is not lost.
What is named can be defended.
What is defended can be reconquered.

NEMO SUPRA LEGEM EST

Amine RAITI · Infrastructure Architect & SRE · Former engineering school professor
Teaching since 2006 · Electricity to Kubernetes · All audiences · CC BY-NC-SA 4.0

وَرَبُّنَا ٱلرَّحْمَٰنُ ٱلْمُسْتَعَانُ عَلَىٰ مَا يَصِفُونَ

Amine RAITI · CC BY-NC-SA 4.0